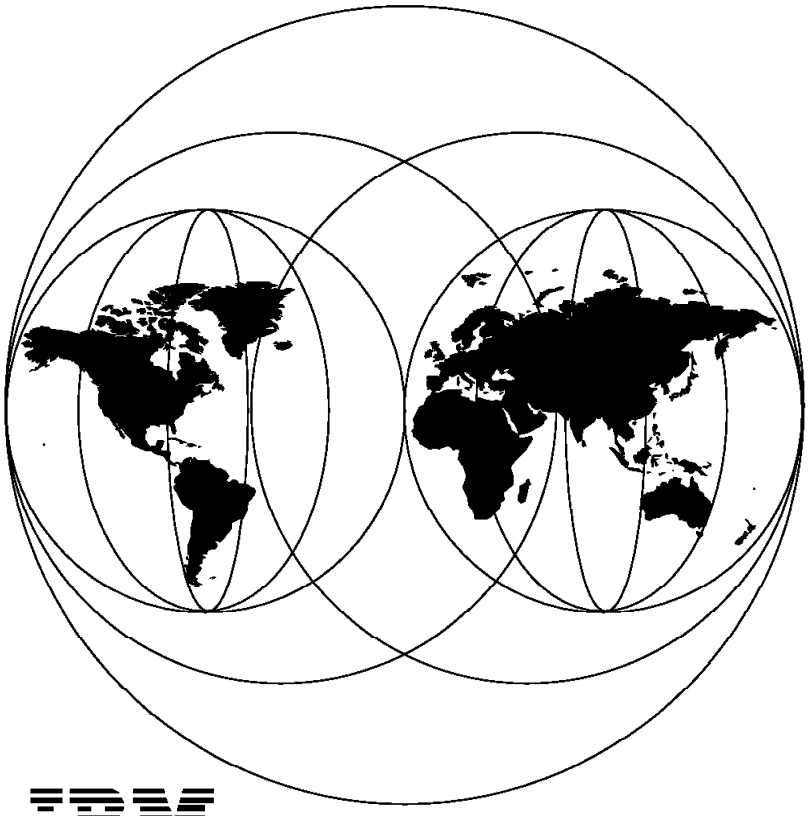
SG24-4824-00

Building the Infrastructure for the Internet

November 1996



IBM

International Technical Support Organization
Raleigh Center



International Technical Support Organization

SG24-4824-00

Building the Infrastructure for the Internet

November 1996

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix H, “Special Notices” on page 615.

First Edition (November 1996)

This edition applies to IBM solutions available in the Internet environment.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. HZ8 Building 678
P.O. Box 12195
Research Triangle Park, NC 27709-2195

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright International Business Machines Corporation 1996. All rights reserved.**
Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

Preface xi
How This Redbook Is Organized xi
The Team That Wrote This Redbook xii
Comments Welcome xiv

Chapter 1. Hardware Platforms 1
1.1 Introduction 1
1.2 Considerations 1
1.2.1 Bandwidth 1
1.2.2 Content Type 2
1.2.3 Number of Clients 2
1.2.4 Servers 2
1.2.5 Scalability 3
1.2.6 Recommendations 4
1.2.7 IBM Servers 6
1.3 Access Technologies 12
1.3.1 Spread Spectrum Technology 12
1.3.2 Leased-Line Connections 16
1.3.3 Cable Modems 17
1.3.4 Integrated Services Digital Network (ISDN) 19

Chapter 2. Networking Hardware 21
2.1 IBM 8235 Dial-In Access to LANs Server 21
2.1.1 8235 System Components 21
2.1.2 Dial-In Access to LANs Server (DIALs) Client Software 22
2.1.3 Using the IBM Dial-Up for TCP/IP 34
2.1.4 IBM 8235 New Features 41
2.1.5 What Is a Virtual Connection? 45
2.1.6 What Is Channel Aggregation? 45
2.1.7 8235 Management Facility 46
2.1.8 8235 Hardware 71
2.1.9 Models Summary 75
2.1.10 Communication Options 76
2.1.11 Supported Protocols 77
2.1.12 Security 86
2.1.13 The Activity Logger 88
2.2 IBM 2210 Nways Multiprotocol Router 89
2.2.1 Models of the IBM 2210 89
2.2.2 Indicators on the IBM 2210 91
2.2.3 The Reset Button on the IBM 2210 92
2.2.4 Networks Supported by the IBM 2210 93
2.2.5 Accessing the IBM 2210 93
2.2.6 Software Package 93
2.2.7 MRNS Overview 94
2.2.8 The IBM 2210 as an IP Router 103
2.2.9 Data Link Switching 114
2.2.10 Features and Facilities 117
2.3 IBM 6611 Router 123
2.3.1 Hardware Overview 124
2.3.2 Multiprotocol Connectivity 129
2.3.3 Bridging with IBM 6611 132

© Copyright IBM Corp. 1996

2.3.4 Data Link Switching . . . 145
2.3.5 IBM 6611 Network Processor Enhancements - Release 4 . . . 153

Chapter 3. Additional IBM Software Solution . . . 155
3.1 Overview . . . 155
3.2 TCP/IP Client/Server Software . . . 157
3.3 Client Software . . . 158
3.3.1 Internet Connection for OS/2 Warp and Windows . . . 158
3.3.2 Warp Connect . . . 159
3.3.3 Secure WebExplorer for OS/2 Warp and AIX . . . 159
3.3.4 WebExplorer for AIX . . . 159
3.4 Server Software . . . 159
3.4.1 Internet Connection Server for OS/2 Warp and AIX . . . 159
3.4.2 Internet Connection Secure Server for OS/2 Warp and AIX . . . 160
3.4.3 Internet Connection Server for MVS/ESA . . . 160
3.4.4 Internet Connection Secure Server for MVS/ESA . . . 160
3.4.5 WebConnection for OS/400 . . . 161
3.4.6 IBM Connection Server Family . . . 161
3.5 Internet Servers . . . 163
3.5.1 Internet POWERsolution for AIX - IBM Internet Connection Servers 163
3.5.2 Internet POWERsolution for AIX - Netscape Servers . . . 163
3.6 Lotus InterNotes . . . 163
3.6.1 Lotus InterNotes Web Publisher . . . 163
3.6.2 Lotus InterNotes News . . . 166
3.7 Other Lotus Software Solutions to the Internet . . . 167
3.7.1 Lotus Domino WebServer . . . 167
3.7.2 Lotus Word Pro . . . 170

Chapter 4. Web Development . . . 175
4.1 Hypertext Markup Language (HTML) . . . 175
4.1.1 HTML2.0 Document Structure . . . 176
4.1.2 HTML2.0 Syntax . . . 176
4.1.3 HTML3.0 or HTML+ . . . 192
4.1.4 HTML Special Symbols . . . 198
4.1.5 HTML Editors and Tools . . . 199
4.1.6 Extensions to HTML . . . 204
4.2 Images . . . 204
4.2.1 HTML Image Files . . . 204
4.2.2 PDF (Portable Document Format) . . . 205
4.2.3 Transparent Images . . . 205
4.3 Other Resources (Audio and Video) . . . 207
4.4 HTML Converters . . . 208
4.4.1 BookMaster to HTML . . . 208
4.4.2 FrameMaker to HTML . . . 212
4.4.3 Interleaf to HTML . . . 213
4.4.4 Other HTML Converters . . . 214
4.5 CGI's Programming . . . 214
4.5.1 The Choice of the Transference Method . . . 215
4.5.2 Catching the REQUEST_METHOD Variable . . . 215
4.5.3 Catching the QUERY_STRING Variable . . . 216
4.5.4 Standard Input on the POST Method . . . 216
4.5.5 Standard Output . . . 216
4.5.6 Decode the Input . . . 217
4.5.7 CGI Variables . . . 218
4.5.8 Content Type considerations . . . 219

4.5.9 Examples, Examples, Examples . . . 219
4.5.10 Ideas for Interesting Pages with CGI Programming . . . 228
4.5.11 Error Handling with CGIs . . . 229
4.5.12 CGI Security Considerations . . . 229
4.6 Virtual Reality Modeling Language . . . 229
4.6.1 VRML specifications . . . 230
4.7 Other Useful Tools . . . 232
4.7.1 JavaScript . . . 233

Chapter 5. Java Programming . . . 235
5.1 Applets and Applications . . . 235
5.1.1 Applets Security Restrictions . . . 236
5.2 Java Basics . . . 236
5.3 Differences between Java and C/C++ . . . 236
5.4 Java Compiler and Interpret . . . 237
5.5 Language Syntax . . . 238
5.6 Variable Types and Declarations on Types . . . 238
5.7 Classes, Objects, Inheritance . . . 242
5.8 Overriding Methods . . . 243
5.9 From Arrays to Loops . . . 244
5.9.1 Casting Elements . . . 244
5.9.2 Conditionals . . . 245
5.9.3 switch... case... default . . . 245
5.9.4 do...while . . . 247
5.9.5 while and for Commands . . . 247
5.9.6 Labeled Loops . . . 248
5.10 Applets Basics . . . 249
5.11 Implementing a Simple Clock . . . 251
5.11.1 The init Method . . . 252
5.11.2 The Start and Stop Methods . . . 252
5.11.3 The Destroy Method . . . 252
5.11.4 The Paint Method . . . 252
5.12 Threading Applets . . . 253
5.13 Graphics on the Applets . . . 255
5.14 Animations, Sounds and Other Effects . . . 258
5.15 Events Handling . . . 266
5.15.1 The Mouse Event Handler . . . 267
5.16 AWT (Abstract Window Toolkit) . . . 268
5.16.1 Layouts and Panels . . . 273
5.17 URL Management . . . 273
5.18 Brief Guide to Advanced Topics . . . 276
5.19 When to Consider CGI and When to Consider JAVA . . . 276

Chapter 6. Multimedia Concepts and Terms . . . 279
6.1.1 JPEG Image Format . . . 279
6.1.2 GIF Image Format . . . 281
6.2 Audio File Formats . . . 283
6.3 Musical Instruments Digital Interface (MIDI) . . . 284
6.3.1 General MIDI Standard . . . 285
6.3.2 MIDI Mapper . . . 285
6.3.3 MIDI Sequencer . . . 285
6.3.4 When to use MIDI . . . 286
6.3.5 Storage Formats . . . 286
6.4 Digital Movie Formats . . . 286
6.4.1 What You Need to Play Movie Files . . . 286

6.4.2 Movie File Formats . . . 287
6.4.3 Movie Players . . . 288
6.5 Multimedia Applications on the Internet . . . 288
6.5.1 IBM Internet Connection Phone . . . 288
6.5.2 Audio on Demand . . . 291
6.5.3 Video Conference . . . 291
6.6 Multimedia Glossary . . . 293

Chapter 7. Existing Gateways . . . 299
7.1 DB2WWW Gateway . . . 299
7.1.1 Installation . . . 300
7.1.2 Configuring DB2WWW . . . 301
7.1.3 The Macro File . . . 301
7.1.4 Accessing Non-DB2 Databases with DB2WWW . . . 309
7.2 Other Database's Gateways . . . 309
7.2.1 Oracle . . . 309
7.2.2 Sybase . . . 309
7.3 MQSeries Gateway . . . 310
7.3.1 Software . . . 310
7.3.2 Installation . . . 310
7.3.3 Gateway Components . . . 311
7.3.4 MQGate . . . 311
7.3.5 Configuration . . . 312
7.3.6 Host Name Sample Application . . . 312
7.3.7 Queue Browser Sample Application . . . 313
7.3.8 CGI Put Sample . . . 314
7.3.9 CGI Get Sample . . . 321
7.3.10 Application Programming Using the Gateway . . . 321
7.3.11 Source Code . . . 322
7.4 AS400 Web Server Screen Translator . . . 324
7.4.1 The 5250 HTML Gateway Server . . . 329
7.4.2 5250 HTML Workstation Gateway Application Logon Exit Program . 330
7.4.3 Configure TCP/IP Workstation Gateway (CFGTCPWSG) Main Menu 331
7.4.4 Change Workstation Gateway Attributes (CHGWSGA) CL Command Prompt . . . 331
7.4.5 How Can I Use the HTML Support for New Possibilities? . . . 335

Chapter 8. Security on the Internet . . . 339
8.1 Policies . . . 340
8.1.1 Organization Issues . . . 340
8.1.2 Who Makes the Policy? . . . 340
8.1.3 Who Is Involved? . . . 340
8.1.4 Responsibilities . . . 340
8.1.5 Risk Assessment . . . 341
8.1.6 Policy Issues . . . 342
8.1.7 Locking In or Out . . . 345
8.2 Establishing Procedures to Prevent Security Problems . . . 346
8.2.1 Identifing Possible Problems . . . 346
8.2.2 Choose Controls to Protect Assets in a Cost-Effective Way . . . 347
8.3 Physical Security . . . 348
8.3.1 Procedures to Recognize Unauthorized Activity . . . 348
8.3.2 Tools for Monitoring the System . . . 348
8.3.3 Vary the Monitoring Schedule . . . 349
8.3.4 Communicating Security Policy . . . 350
8.3.5 Resources to Prevent Security Breaches . . . 352

8.3.6 Problem Reporting . . . 369
8.3.7 Secure Web Servers . . . 369
8.3.8 IBM Internet Connection Secure Products . . . 373
8.3.9 Eletronic Commerce . . . 374
8.3.10 Reference Sites on the Internet . . . 383

Chapter 9. Network Management . . . 385
9.1 SystemView Introduction . . . 385
9.1.1 SystemView Benefits . . . 385
9.1.2 SystemView Structure . . . 386
9.2 Managing a Heterogeneous Network . . . 388
9.2.1 A Brief View into SNMP History . . . 388
9.2.2 SNMP Definitions . . . 390
9.2.3 The SNMP Architecture . . . 391
9.2.4 Goals of the SNMP Architecture . . . 392
9.2.5 SNMP Model . . . 392
9.2.6 User Datagram Protocol (UDP) . . . 393
9.2.7 Asynchronous Request/Response Protocol . . . 394
9.2.8 SNMP Agent . . . 395
9.2.9 SNMP Subagent . . . 395
9.2.10 SNMP Manager . . . 396
9.2.11 SNMP Version 2 . . . 397
9.2.12 Understanding MIBs . . . 399
9.2.13 SNMP Operations . . . 399
9.2.14 Desktop Management Interface (DMI) . . . 400
9.3 Overview of IBM Products for Network Management . . . 402
9.3.1 Positioning the AIX Management Platform . . . 403
9.3.2 AIX Management Platform Overview . . . 405
9.3.3 Positioning the MVS Management Platform . . . 406
9.3.4 MVS Management Platform Overview . . . 408
9.3.5 Positioning the OS/2 Management Platform . . . 410
9.3.6 OS/2 Management Platform Overview . . . 412
9.3.7 Positioning the Windows IBM Management Platform . . . 413
9.3.8 Windows IBM Management Platform Overview . . . 414
9.3.9 Tivoli TME 10 . . . 415
9.4 More Information . . . 416

Chapter 10. Connection Access Services . . . 419
10.1 Service Providers . . . 419
10.1.1 How to Select an Internet Service Provider . . . 419
10.1.2 How to Build an Infrastructure for an Internet Service Provider . . . 422
10.2 IBM As a Service Provider . . . 426
10.2.1 IBM Global Network . . . 426
10.2.2 Advantis . . . 427
10.2.3 Prodigy Services Company . . . 427
10.3 IBM Internet Connection Access Services . . . 428
10.3.1 Dial-Up Services . . . 428
10.3.2 Corporate Dial Services . . . 432
10.3.3 Leased Line Internet Connection Services . . . 434
10.3.4 IGN's Internet Backbone Design . . . 441

Chapter 11. Content Services on the Internet . . . 451
11.1 The Basic Internet Services . . . 451
11.1.1 The World Wide Web . . . 451
11.1.2 Web Farms Concept . . . 452

11.1.3 Communication Services . . . 452
11.1.4 Information Search and Retrieval Services . . . 452
11.2 Content Services Concept . . . 453
11.3 Content Services through the IBM Global Network . . . 453
11.3.1 Highlights . . . 453
11.3.2 Enhanced Services . . . 453
11.3.3 Versatility and Security . . . 454
11.3.4 Priced for Performance . . . 454
11.3.5 Operating Environment . . . 454
11.3.6 Connectivity to the Internet . . . 454
11.3.7 IBM Domain Name Services . . . 454
11.3.8 Monthly Server Activity Report . . . 455
11.4 Creating a Content Hosting Service . . . 455
11.4.1 Content Hosting Description . . . 455
11.4.2 Hardware Requirements . . . 455
11.4.3 Software Requirements . . . 456
11.4.4 Connection Requirements . . . 456
11.4.5 Network Solution Design Sample . . . 457
11.4.6 IP Addressing . . . 457
11.4.7 Domain Name Systems . . . 459
11.4.8 The Flat Name Space . . . 460
11.4.9 The Domain Name System . . . 460
11.4.10 Mail Support . . . 479
11.4.11 DNS Design Requirements . . . 479
11.4.12 Web Server Softwares . . . 484
11.4.13 Multiple IP Addresses . . . 485
11.4.14 Setting up the Netscape Commerce Server for AIX . . . 485
11.4.15 Running the Netscape Commerce Server Administration . . . 507
11.4.16 Putting Web Content on the Internet . . . 508
11.4.17 Working with CGI programs . . . 509
11.4.18 Developing an FTP Site . . . 512
11.4.19 Getting Reports from Content Services . . . 516
11.4.20 Network Monitoring and Management . . . 518
11.4.21 Network Management . . . 519
11.4.22 Operational Issues . . . 520
11.4.23 Security . . . 520

Chapter 12. Networked Applications . . . 523
12.1 IBM infoMarket . . . 523
12.1.1 Wide Area Search of Distributed Data . . . 524
12.1.2 Cryptolope . . . 526
12.1.3 Cryptolope Container . . . 526
12.1.4 Key IBM infoMarket Rights Management Directions . . . 526
12.1.5 Superdistribution . . . 527
12.1.6 IBM infoMarket Applications . . . 529
12.2 IBM infoSage . . . 531
12.2.1 Profile Editor . . . 532
12.2.2 Special Editions . . . 533
12.2.3 Archive Search . . . 533
12.2.4 Stock Tracker . . . 533
12.2.5 Top Stories . . . 533
12.2.6 Links . . . 533
12.3 Electronic Purchasing Service . . . 534
12.4 Interactive Marketing Service . . . 536
12.5 Net.Commerce . . . 537

12.5.1 The Store Manager . . . 538
12.5.2 The Store Creator . . . 538
12.5.3 The Store Administrator . . . 539
12.5.4 The Template Editor . . . 539
12.5.5 The Net.Commerce Director . . . 539
12.5.6 The Net.Commerce Daemon . . . 539
12.5.7 The Lotus Payment Switch . . . 540
12.5.8 The Olympic Ticket Sale - an Example of Net.Commerce . . . 540

Chapter 13. Internet Sample Solutions . . . 545
13.1 Basic E-mail Solution . . . 545
13.2 Corporative Secure LAN Solution . . . 547
13.3 Electronic Commerce Solution . . . 548

Chapter 14. Consulting Services . . . 553
14.1 Management Information Technology Consulting Service Lines . . . 553
14.1.1 Transformation Services . . . 553
14.1.2 I/T Consulting Services . . . 554
14.1.3 Integration Services . . . 554
14.1.4 General Business . . . 555
14.1.5 Technology Services . . . 555
14.2 Industry Specializations . . . 555
14.3 Internet Consulting and Services . . . 556

Appendix A. The IAB . . . 559
A.1 The Internet Activities Board (IAB) . . . 559
A.1.1 Request for Comments (RFC) . . . 560
A.1.2 Functions of the IAB . . . 561
A.1.3 Protocol Standardization Process . . . 561

Appendix B. A Brief Description of IBM Network Management Products . . . 565
B.1 AIX Platform . . . 565
B.2 MVS Platform . . . 568
B.2.1 Basic Products . . . 568
B.2.2 Optional Products . . . 569
B.3 OS/2 Platform . . . 573
B.4 Windows Platform . . . 575

Appendix C. IBM infoMarket Rights Management Architecture . . . 577

Appendix D. More Information about IBM infoSage . . . 581
D.1 Content Resources . . . 581
D.2 IBM Profile Editor Screens . . . 584
D.3 IBM infoSage Result Examples . . . 588
D.3.1 Morning Delivery . . . 588
D.3.2 Special Edition . . . 589
D.3.3 Alert Delivery . . . 592

Appendix E. IBM Global Network Phone List . . . 595

Appendix F. IBM Global Network Registration Phone List . . . 611

Appendix G. IBM Global Network Help Desk Phone List . . . 613

Appendix H. Special Notices . . . 615

Appendix I. Related Publications 621
I.1 International Technical Support Organization Publications 621
I.2 Redbooks on CD-ROMs 621
I.3 Other Publications 622

How To Get ITSO Redbooks 623
How IBM Employees Can Get ITSO Redbooks 623
How Customers Can Get ITSO Redbooks 624
IBM Redbook Order Form 625

Index 627

Preface

This redbook provides detailed coverage of the Internet, focusing on solutions available for the Internet environment. It includes information on hardware (remote connection, routers, and servers), software (client, servers, browsers, and TCP/IP), and services available to build an Internet infrastructure in any company. It also addresses management systems, gateways to databases, and the application development environment, with some details on the language being used. This book also provides information on the booming Web environment and how to access it, and how to develop Web pages, including details on the HTML and JAVA languages and the integration of such multimedia elements as video and audio.

This redbook provides readers with a broad view of all solutions available in the Internet environment, helping them to select the solution that is most suitable for their companies′ needs. The detailed descriptions of the services is very important for readers to decide how to ″put their best food forward″ on the Internet.

This redbook was written for customers, IBM technical professionals, service specialists, marketing specialists and marketing representatives working in the Internet area.

Some knowledge of networking and the application environment is assumed.

How This Redbook Is Organized

This redbook contains 644 pages. It is organized as follows:

- Chapter 1, “Hardware Platforms”

 This chapter provides a description of the basic hardware available to build your Internet servers and clients and the access technologies available.

- Chapter 2, “Networking Hardware”

 This chapter discusses the networking equipment available to connect your environment to the network including routers and remote access.

- Chapter 3, “Additional IBM Software Solution”

 This chapter provides the software in the Internet environment provided by IBM.

- Chapter 4, “Web Development”

 This chapter discusses the technology available to develop Web pages in the Internet.

- Chapter 5, “ Java Programming”

 This chapter provides information on the Java programming including samples.

- Chapter 6, “Multimedia Concepts and Terms”

 This chapter discusses the concepts of multimedia technology used in the Internet environment.

- Chapter 7, “Existing Gateways”

© Copyright IBM Corp. 1996

This chapter provides the gateway solutions by IBM to the Internet including DB2, MQSeries and CICS to the Web environment.

- Chapter 8, “Security on the Internet”

 This chapter discusses descriptions on security policies and procedures available for the Internet environment including secure servers and secure electronic transactions.

- Chapter 9, “Network Management”

 This chapter provides the basic concepts of network management applicable to the Internet and the description of products available in different operating system environments.

- Chapter 10, “Connection Access Services”

 This chapter discusses the description of an Internet service provider and the connection services offered by the IBM Global Network.

- Chapter 11, “Content Services on the Internet”

 This chapter provides the description of content services and how it is provided by the IBM Global Network.

- Chapter 12, “Networked Applications”

 This chapter discusses some different applications available in the Internet such as InfoMarket, InfoSage and NetCommerce.

- Chapter 13, “Internet Sample Solutions”

 This chapter provides some basic Internet solutions such as e-mail, eletronic commerce and secure LAN.

- Chapter 14, “Consulting Services”

 This chapter discusses information on available IBM consulting services.

The Team That Wrote This Redbook

This redbook was produced by a team of specialists from around the world working at the Systems Management and Networking ITSO Center, Raleigh.

Ricardo Haragutchi is a senior ITSO specialist at the Systems Management and Networking ITSO Center, Raleigh. He writes extensively and teaches IBM classes worldwide on all areas of LAN hardware and the Internet environment. Before joining the ITSO one year ago, he worked in the Field Systems Center (FSC), IBM Brazil as a senior system engineer.

Barry D. Nusbaum is a senior ITSO specialist for AIX, OS/2 and NT Systems Management at the Systems Management and Networking ITSO Center, Raleigh.

Carlos de Luna Sáenz is an information technology specialist in Mexico. He has two years of experience in Internet and Web developing fields. He holds a degree in Computer Science Engineering from the Instituto Tecnológico y de Estudios Superiores de Monterrey, Campus Estado de México. His areas of expertise include CGI and Java programs for Internet developing and database access.

Nilson Tenorio Batista is a system support specialist in Brazil. He has three years of experience in information system security, four years of experience in networking technology and two years of experience in Internet content services.

He has worked at IBM Brazil for six years and two years for IBM Global Network. He holds a degree in Data Processing Technology from the Pontificia Universidade Catolica do Rio de Janeiro - PUC/RJ. His areas of expertise include Internet content services development and support.

Roberto Morizi Oku is a system support specialist in Brazil. He has three years of experience in networking technology. He has worked at IBM/GSI Brazil for five years. His areas of expertise include IBM Internet Connection Access Services, IGN Dial Services, and INGW (Intelligent Network Gateway) and LIG (Local Gateway Interface) technical support. He holds a degree in Engineering from the Escola Politécnica, University of São Paulo, Brazil.

Patrick Schmitt-Heinrich is a network systems specialist in Germany. He has six years of experience in the IBM Global Network. He holds a degree in telecommunications from the Staatliche Studienakademie Baden-Wuerttemberg. His areas of expertise include the design of network solutions in the Internet area.

Robert Macgregor is a technical support specialist at the Systems Management and Networking ITSO Center, Raleigh, dealing with open systems management and network security topics. Under his technical leadership, 10 redbooks have been published, including books on the Internet Connection Secure Network Gateway, NetView for AIX and SystemView for AIX products. Before coming to the ITSO, Rob provided technical support, services and consultancy for IBM customers in the United Kingdom.

Thanks to the following people for their invaluable contributions to this project:

Alfred Christensen
David L. Boone
Cameron Ferstat
David Shogren
Guido de Simoni
Mark Kressin
Sergio Juri
Silvio Podcameni
Roger Serrette
International Technical Support Organization

Diana Reese
Phillip A. Rzewski
Sarah Siegel
Shawn Tien
Advantis

Adriana F. Lira
Arthur Ryman
Charles Poland
Gene Tsudik
John Sweeney
Marcela Toledo
Mark H. Linehan
Sal A. Calta
IBM

Paul Holbrook
CICNet Inc.

Marcus J. Ranum
V-ONE Corporation

Brent Chapman
Great Circle Association

Brian Boyle
Exxon Research

Comments Welcome

We want our redbooks to be as helpful as possible. Should you have any comments about this or other redbooks, please send us a note at the following address:

redbook@vnet.ibm.com

Your comments are important to us!

Chapter 1. Hardware Platforms

This chapter contains useful information about all IBM platforms that can be used as Internet servers, describes processor technologies, operating systems, adapters, and gives you the necessary data to do efficient server capacity planning.

1.1 Introduction

The Internet has been growing at a phenomenal pace, connecting each new user with a vast amount of global information covering every interest from classic cars to politics to investments. Organizations put their Web servers on the Internet to make their products and information more accessible to a global audience.

Sizing a Web server for the Internet can be a very difficult task. The Internet includes millions of interconnected individuals who are navigating from one Web server to the next in search of information that has value to them.

Rapid advances in Internet technology are changing the way we work. New technologies of software and hardware are announced every day. Selecting the proper server hardware is vital to those who want to be productive now and in the future. Internet applications need servers capable of providing information that is available full-time with good performance.

Availability and performance are fundamental requirements when we talk about servers that will be connected on the Internet. There is no Internet user that likes to wait to receive information. You need to guarantee that your server will deliver information faster so that these users will want to be consumers of your products and services.

Today you can use all existing platforms to deliver information on the Internet, such as Intel and RISC-based machines, AS/400 and mainframes. You need to choose the system that fills your performance needs and investment limits.

1.2 Considerations

The following sections describe the considerations necessary when choosing a hardware system.

1.2.1 Bandwidth

In working with a customer to size up a Web solution, it is important to understand the implications of the speed of the networking connection to the Web server. More often than not, many potential Web content providers are very focused on the vague *hits per day* quantity. The level of traffic that a particular Web server can support will be dependent on the server type, the content accessed on the server and the speed of the connection of the server to the intra/Internet environment.

An Internet service provider will deliver a connection of defined speed; five of these most common speeds are: leased lines between 56 Kbps and 256 Kbps ISDN (128 Kbps), T1 (1.544 Mbps), and T3 (45 Mbps). For an intranet environment,

© Copyright IBM Corp. 1996

common LAN speeds are 10 Mbps (over Ethernet), 16 Mbps (over token-ring) and 100 Mbps (over fast Ethernet or FDDI).

As the average Web transaction size increases, the maximum number of transactions decreases. Sites that plan on being mostly text-based will have average transactions sizes around 1 to 5 KB; most well-designed sites with a mix of text and graphics intended for access by modem users are in the 10 KB per transaction size and sites with a substantial portion of multimedia content can exceed 100 KB per transaction.

1.2.2 Content Type

The physical size of the Web content is important in looking at the resources required for a server, indicating the necessary data storage requirements. Additionally, when the content on the Web server is dynamically generated, substantial processing resources may be required. Dynamic content on a Web site can be generated in many ways, from a simple counter that displays the number of hits that a page has received, to a system that uses analysis of user clicks to tailor the information (and advertisements in some cases) that the user sees at the site.

1.2.3 Number of Clients

The number of simultaneous users of a site is very challenging to characterize. Unlike other types of client/server architectures, the weight of an individual client on the Web server is quite small and short-lived. Connections to a Web server are traditionally stateless sessions that begin with an open from the client, a request for data, a server reply with data, and then the session closes. Depending on the speed of the network connection, the size of the data requested and the server load, this session can last from tenths to tens of seconds.

A major portion of the content on the Web is static. This includes both images and textual data. The CPU resources required to serve such data are minimal. The IBM server products have a large performance range from basic Intel processor-based systems to highly parallel processing servers.

A typical http connection consists of a client open, client request, server header and data response and connection shutdown. The average response size is approximately 7 KB.

When a Web server responds to users in a more dynamic way, we see a much stronger case for increased computing power at the server. In some configurations, there are still situations where the performance is network bound.

1.2.4 Servers

You need to choose the perfect combination between a hardware platform and the operating system. This is because some platforms do not support the newest powerful applications that can be useful to improve the quality of your Internet server.

Some companies use an existing operational platform as the Internet server. It can be a problem if this server has confidential documents, corporative applications and highly secure data. A hacker will be able to steal or destroy this important data using daemons such as HTTP, GOPHER, and FTP servers as

gates to go inside your system. The best option is to create a server on a dedicated machine that will be exposed to the Internet without any confidential data. The majority of servers connected to the Internet are running on UNIX systems on RISC-based machines, but today a lot of new servers running OS/2, WindowsNT and Linux on Intel-based machines are being used. Some companies are also using mainframes running VM and MVS and AS/400 as servers. The following table shows the available services on each platform.

Table 1. Available Services on Different Operating Systems

Operating System	DNS	E-mail	GOPHER	HTTP	TELNET	FTP	NEWS	DB/2	LOTUS NOTES	JAVA
AIX	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
OS/2	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
NT	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
OS/400	NO	YES	YES	YES	YES	YES	NO	YES	YES	NO
MVS	YES	YES	YES	YES	YES	YES	NO	YES	NO	NO

1.2.5 Scalability

The demand for scalable systems is growing. Stated simply, a scalable system is one that permits the addition of processing power, storage, memory, input/output (I/O), and connectivity with relative ease, so user organizations can deploy larger, more complex, more sophisticated applications to exploit constantly growing databases and make both available to increasing numbers of users through very high bandwidth networks.

Technically, the simplest way to provide scalability is to build larger and faster uniprocessors. Systems can also be made faster using highly sophisticated architectures (either alone or in combination with unique technologies). The advantage of scaling uniprocessors is that the software remains the same; it simply runs on a faster processor.

One can also scale by integrating multiple uniprocessors into a single system in which they share resources such as memory, I/O, the operating system, and application software. Having one of each resource makes a symmetric multiprocessor (SMP) system relatively easy to program and manage. In addition, the SMP will run essentially the same software as the uniprocessor, although it may have to be modified to remove bottlenecks that the faster multiprocessor could expose.

Another way to get scalability is to use parallel systems where multiple processors are connected to each other by a high-performance interconnect mechanism. Each processor has its own memory, its own I/O configuration, and its own copy of the operating system. Thus, far higher levels of scalability are achievable. Indeed, such systems become almost infinitely scalable because the incremental processor does not increase contention for resources; it comes with all it needs to do productive work.

The AIX systems can scale efficiently to four or eight processors using PowerPC technology on SMP systems. So, using parallel systems based on Power and Power2 processors, AIX can deliver extremely high performances. Because it′s relatively new, NT does not scale nearly as well as UNIX. Theoretically, NT is designed to support up to 32 processors; in reality it is currently limited to four

processors in most situations. Depending on the mix of applications and hardware architectures, the number of processors can be as low as two or as high as eight. The OS/2 can scale up to 16 processors on the Warp Server version and is a good choice for Internet applications that demand performance and integration with CICS, IMS and DB/2. If you're writing in-house applications for multiprocessor systems, you must write code so that instructions are handled as a series of threads. This lets the operating system efficiently direct processes to different CPUs.

Table 2. Operating System and Minimum Configuration to a Basic Web Server

Operating System	**Recommended minimum configuration**
AIX	• IBM RS/6000 - Model 43P - 100 MHz CPU • RAM - 64 MB • Hard disk - 2.0 GB • CD-ROM • 15″ Display • AIX 4.1.4
OS/2 and WindowsNT	• IBM PC Server 310 - Pentium 100 MHz CPU • RAM - 32 MB • Hard disk - 2.0 GB • CD-ROM • 14″ Display • PCI or EISA Ethernet adapter • OS/2 Warp Server or WindowsNT 3.5.1
OS/400	• IBM AS/400 - Model 20S - 64-bit PowerPC CPU • RAM - 48 MB • Hard disk - 3.0 GB • Tape drive • 5250 console display • Ethernet adapter • OS/400 V3.R6
MVS	• Any S/390 • MVS Operating System • TCP/IP for MVS • IP connection using a LAN or WAN

1.2.6 Recommendations

The basic Internet structure is the World Wide Web (WWW) server and the e-mail server. You can use other resources such as the FTP server, Telnet server, Database server, Gopher server, News server, Chat server, and DNS server, but the WWW server and the e-mail server are all you need to create an initial Internet structure. Depending on the hardware technology and the power of your server, you can run some of these server daemons on same machine. When the performance needs to increase, you will need to improve server performance or divide these daemons on other servers.

Creating an Internet structure can be a low, medium or high-cost investment; it depends on the type of service and information that you will provide on the Internet. In general, Internet sites that are connected by T1 lines and Ethernet-LAN connected intranet sites with largely static data, are adequately served by a entry uniprocessor system with adequate disk storage for the content provided. It is important to have enough RAM to accommodate both the http server processes and for file caching of page content that resides on disk. Sites with high-bandwidth connections to the Internet and intranet sites that can utilize FDDI will benefit from mid-range and SMP solutions. Sites that will generate significant Web content in response to user actions or potential E-Commerce sites should consider such systems even if they are connected by T1 lines to the Internet or Ethernet-LAN to the intranet.

Table 3. How to Calculate Maximum HTTP Operation/Sec for a Determinated Bandwidth and File Size

Network connection type	**Bandwidth kbps**	**File average size - 1 KB**	**File average size - 10 KB**	**File average size - 100 KB**
9.6 modem	9.6 kb	1.2	0.1	0.0
14.4 modem	14.4 kb	1.8	0.2	0.0
28.8 modem	28.8 kb	3.6	0.3	0.0
56 kb leased	56 kb	7.0	0.7	0.1
64 kb leased	64 kb	8.0	0.8	0.1
ISDN 1	64 kb	8.0	0.8	0.1
ISDN 2	128 kb	16.0	1.6	0.2
T1	1.5 Mb	187.5	18.7	1.8
Ethernet	10 Mb	1250.0	125.0	12.5
T3	45 Mb	5625.0	562.0	56.2
FDDI	100 Mb	12500.0	1250.0	125
Fast Ethernet	100 Mb	12500.0	1250.0	125
ATM/155	155 Mb	19375.0	1937.0	193.0
ATM/622	622 Mb	77750.0	7775.0	777.0

Using the values listed on the table above, we can create a hypothetical example of how to conduct Web server capacity planning:

Consider a site with the following characteristics:

- Bandwidth/user.....(2.5 kbps)/user (modem users)
- Average file size.....7 KB/operation
- (operations/sec)/user.....0.35(operations/sec)/user
- # of active users connected.....100 users
- 20% of the users are active at any given time.....20 users
- Then the requirements are:
- Bandwidth.....49 kbps = approx. 1 ISDN1 + 1 ISDN2 Channels
- Operations/sec.....7
- Minimum network sub-system required.....10 Mbps Ethernet

Suppose system X can do 10 operations/sec, then you need only one.

Table 4 shows the questions that can help you choose the right platform to fit your needs.

Table 4. Main Questions to Consider before Configuring a Server

Questions	Commentary
Should AIX, OS/2, VM or Windows NT serve as the Internet server platform?	You need to consider your budget, people skills, your existing in-house environment and performance needs before choosing one platform.
How many hits per day on the server?	You can use this information to do an effective capacity planning. Generally, on a low-hit site you can use an Intel platform, and on a high-hit site it is indicated that you use RISC-based machines.
What are the pages medium size?	You can multiply the medium page size (KB) by the number of hits daily on the server and obtain how much information will be delivered.
Must your external users have access to the databases?	If yes, you will need a more powerful server because in most cases the database gateway daemon degenerates the system performance.
If so, what type of database support is required, such as IBM DB/2, Oracle, Sybase, Ingress or Informix integration?	The database gateways can have different behaviors. First contact your database supplier to check the needs of this software.
What are your security requirements? For example, will it be necessary to protect highly confidential information and restrict access to the internal corporate network?	If yes, you will need a secure server that supports SSL or S-HTTP. This server gets part of the processor power to make security validations.
Will multiple home pages be installed on the same server?	If yes, first consider all the questions listed above, and if necessary add additional memory and/or processor power on your server.
What type of interface do you need to use? It must be intuitive, Motif or Windows-like and easy to use?	This is a very important item when you do not have specialized skills on different platforms. The Windows and Motif-based operating systems such as WindowsNT, AIX X-Windows and OS/2 are easier to use, administrate and install. The VM, MVS and OS/400 operating systems do not support graphical applications.

1.2.7 IBM Servers

IBM can provide Internet solutions on any hardware platform. Here you can see the main products available on each technology that fill the requirements in performance and reliability to be an Internet server.

1.2.7.1 IBM PC Server Family

PC Servers are a good choice for a wide range of Internet applications, creating a scalar and low-cost solution. You can initialize your Internet site using a PC Server with basic features and, depending on the model that you choose, improve the processor power, memory, storage and communication capability. There are a lot of operating systems available to the Intel platform that can perform an Internet server solution. They are as follows:

- IBM OS/2 Warp Connect
- IBM OS/2 Warp Server

- Microsoft Windows 3.1
- Microsoft Windows95
- Microsoft WindowsNT Family
- SCO UNIX
- Linux
- Solaris
- Novell NetWare
- Novell UnixWare

There is an available solution from IBM that is a bundled hardware and software kit including IBM PC Server 320, 32-MB RAM, 2.25-GB hard disk, CD-ROM, operating system, Internet server software, end-user documentation and integrator documentation which comprise a ready-to-build solution for quick installation.

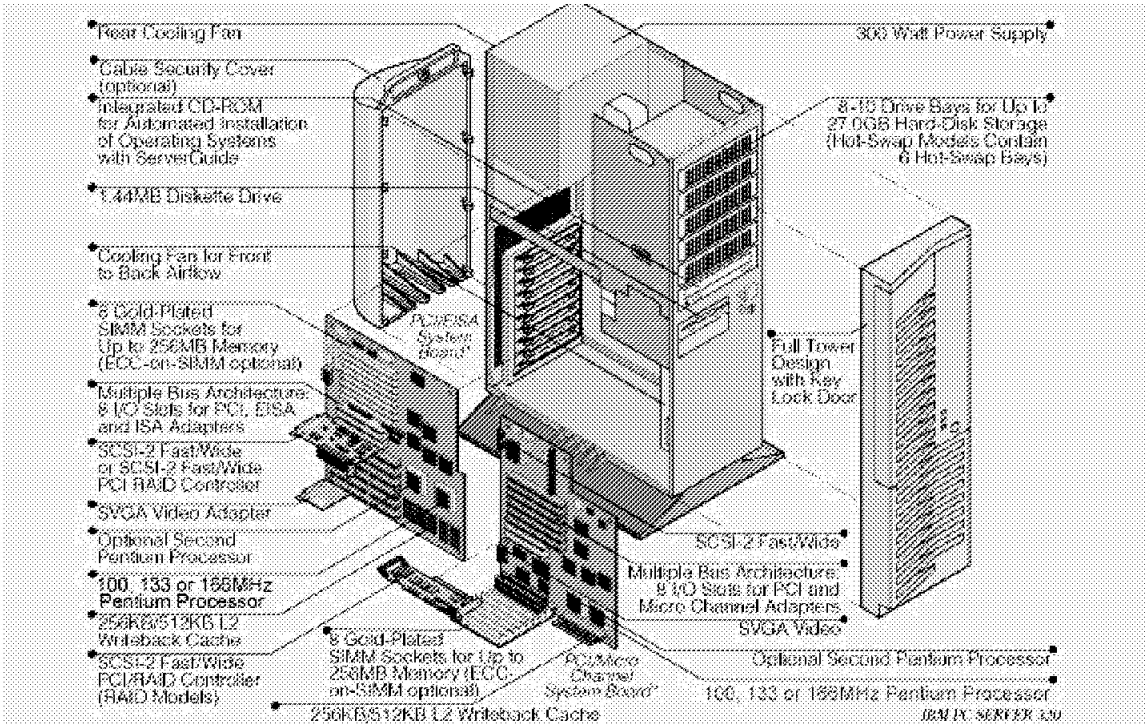


Figure 1. IBM PC Server 320

The available operating system and server choices are:

- IBM OS\2 Warp Server and IBM Internet Connection Secure Server
- Microsoft Windows NT Server and Netscape Commerce Server
- SunSoft Solaris and Netscape Commerce Server

The secured commercial Web server software from IBM or Netscape is included in the kit. The Web presence you create with the PC Server Internet Series will be able to handle queries from Internet users anywhere in the world via industry-standard browsers such as IBM OS/2 Web Explorer and Netscape

Navigator. The OS/2 package allows HTML browser access to CICS and DB2 applications.

If you need more information such as available models, supported devices and technical details about the IBM PC Server family go to the IBM Personal Computing home page on the Internet at http://www.pc.ibm.com.

1.2.7.2 IBM RS/6000 Family

RS/6000 servers are powerful, cost-effective systems with excellent growth and availability options to meet the needs of network-based applications such as the Internet server, Notes server and database server. Customer investment is protected when the new future RS/6000 technologies become available.

IBM′s Internet RS/6000 solutions contain the hardware and software that you need to establish your presence on the Internet. These solutions are designed to operate in a multivendor, networking environment.

The IBM AIX implementation of Sun′s Java programming environment (AIX 4.2 only) helps you deliver your Web page content in a more visually compelling way. For example, it allows you to easily add multimedia and create applications that will be accessed worldwide using the Internet. So, you can have a choice of AIX Web servers available from IBM and Netscape.

One of the main advantages of IBM′s Internet offerings is that you get the power and versatility of UNIX in communications, connectivity, and broad range of optional systems management tools without having an in-depth knowledge of UNIX. Another advantage is the scalability of POWER, POWER2 and PowerPC technologies. From entry servers to parallel systems, RS/6000 can deliver scalar levels of performance.

IBM′s family of Internet POWERsolutions for AIX contain factory-tested and pre-installed hardware and software to establish your presence and conduct business on the Internet′s World Wide Web.

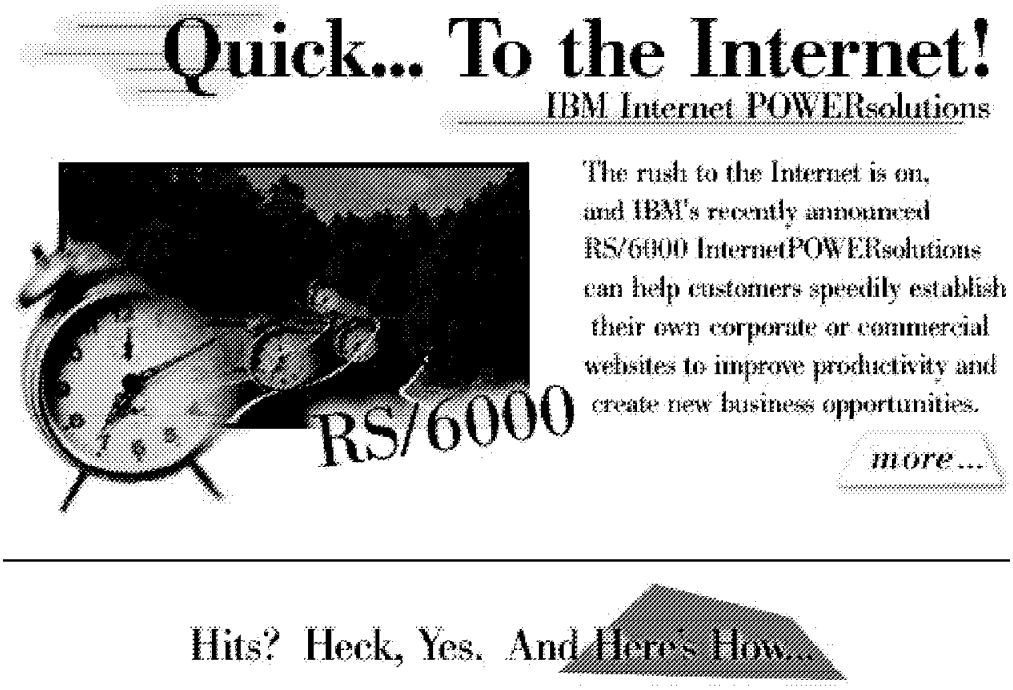


Figure 2. IBM RS/6000

With these Internet POWERsolutions, you can be up and running relatively quickly on the Web. A few steps and your customers or employees are ready to surf. The solution is designed to operate in multivendor, networking environments.

You can choose a solution that contains:

- Secure Web servers for both Internet and intranet needs
- Firewall software for a secure interface between an internal network and the Internet
- Proxy services software for replicating Web page content locally
- Commercial applications for quickly and cost-effectively establishing a full-scale commerce Web site

The solutions take advantage of the scalable capacity of RS/6000 systems, from desktop clients and servers to symmetric multiprocessors to high-powered rack-mounted servers and scalable POWERparallel systems.

Internet software choices can be:

- IBM′s Internet Connection Secure Server
- Netscape′s FastTrack, Enterprise, and Proxy Servers
- IBM′s Internet Connection Secured Network Gateway (firewall)
- Netscape′s Publishing and Community Systems commercial applications

All systems are preconfigured, pretested, and integrated. With an additional option, you can integrate existing business applications, such as DB2 databases and CICS transaction systems on the HTML pages.

The integrated IBM AIX implementation of Sun's Java programming environment (not available with the firewall server) can help deliver Web page content in a more visually compelling way, such as adding animation. A main advantage of IBM's Internet POWERsolution offerings is that you get the power and versatility of UNIX (communications, connectivity, broad range of optional systems management tools, and sophisticated middleware) without having an in-depth knowledge of the operating system.

An Internet POWERsolution with Netscape Proxy Server offers a high-performance solution for replicating and filtering access to Web page content transparently to end users. Requests for specific Web pages are automatically routed to the proxy server, which provides the pages from its local cache. You can even download a group of Web pages and make them locally available. This efficient resource usage can help reduce network costs while giving users a fast, timely response.

These Internet POWERsolutions are backed by IBM's worldwide on-site service and support.

If you need more information such as available models, supported devices and technical details about the IBM RS/6000 family go to the IBM RS/6000 home page on the Internet at http://www.austin.ibm.com.

1.2.7.3 IBM AS/400 Family

The AS/400 platform is an excellent choice to create an Internet server because Internet Connection for AS/400 supports HTTP drivers that can serve any native AS/400 application without a rewrite or recompile over the Internet. Even traditional, host-based applications can be served to terminals running popular Web browsers. Internet users are also able to download files or software, as well as access the AS/400 database, from Web browsers.

Using the HTTP protocol, customers can enhance existing AS/400 applications with hypertext capabilities or attention-getting graphics, audio and video. With Internet Connection, users can also monitor the attention people are paying to their presences on the Web.

AS/400 supports the TCP/IP Serial Link Internet Protocol (SLIP), which provides native TCP/IP connectivity to the Internet over telephone lines.

AS/400 also supports the popular Internet Post Office Protocol (POP3), enabling AS/400 to deliver electronic correspondence to OS/2, UNIX, Windows and Macintosh clients running the most popular mail products.

With support for Lotus Notes Release 4, AS/400 users can use a solution that integrates messaging, groupware and the World Wide Web for building and distributing custom client/server, Internet and intranet applications.

Notes open architecture leverages and maximizes existing AS/400 investments by providing a client/server application development environment, bidirectional field-level replication, client/server messaging and integration with relational databases. Lotus Notes also provides Internet integration, allowing users to publish, locate and share Internet information through functions included in Notes Release 4. Lotus Notes will reside under OS/2 on a dedicated AS/400 Integrated PC Server (FSIOP). The Integrated PC Server can manage up to eight networks, consisting, for example, of Notes, OS/2 or Novell NetWare.

AS/400 has an integrated operating system that provides unrivaled security on the Internet. AS/400 security features protect against hackers and viruses.

If you need more information such as available models, supported devices and technical details about AS/400 Family go to the IBM AS/400 home page on the Internet at http://www.as400.ibm.com.

1.2.7.4 IBM System/390

With System/390, you can meet the needs of thousands of Internet and intranet users. As a server designed for large-volume transactions, it can easily handle just about anything in global networking.

System/390 lets you link existing applications to the World Wide Web with minimal modifications and without moving data to other Web-serving platforms. The IBM Internet Connection Server for MVS/ESA has a direct connection to CICS, IMS, DB2 and MQSeries. The System/390 allows you to start small on your Internet and intranet offerings, then scale up as needed to handle thousands of transactions.

The System/390 can rely on cryptography functions to protect your data. You can establish a wide range of security measures and procedures, such as access control policies, passwords, and special user privileges.

Built into the current Internet Connection Server for MVS/ESA, through the System Access Facility, is access to such MVS system resource managers as RACF or the OS/390 security server. You can use this technology to control access to files and other system resources.

Instead of adding servers to meet changing performance demands, you can allocate System/390 server capacity to the public network partition.

So, System/390 gives you all the security and performance that you need to create a powerful Internet server.

If you need more information such as available models, supported devices and technical details about System/390 go to the IBM System/390 home page on the Internet at http://www.s390.ibm.com.

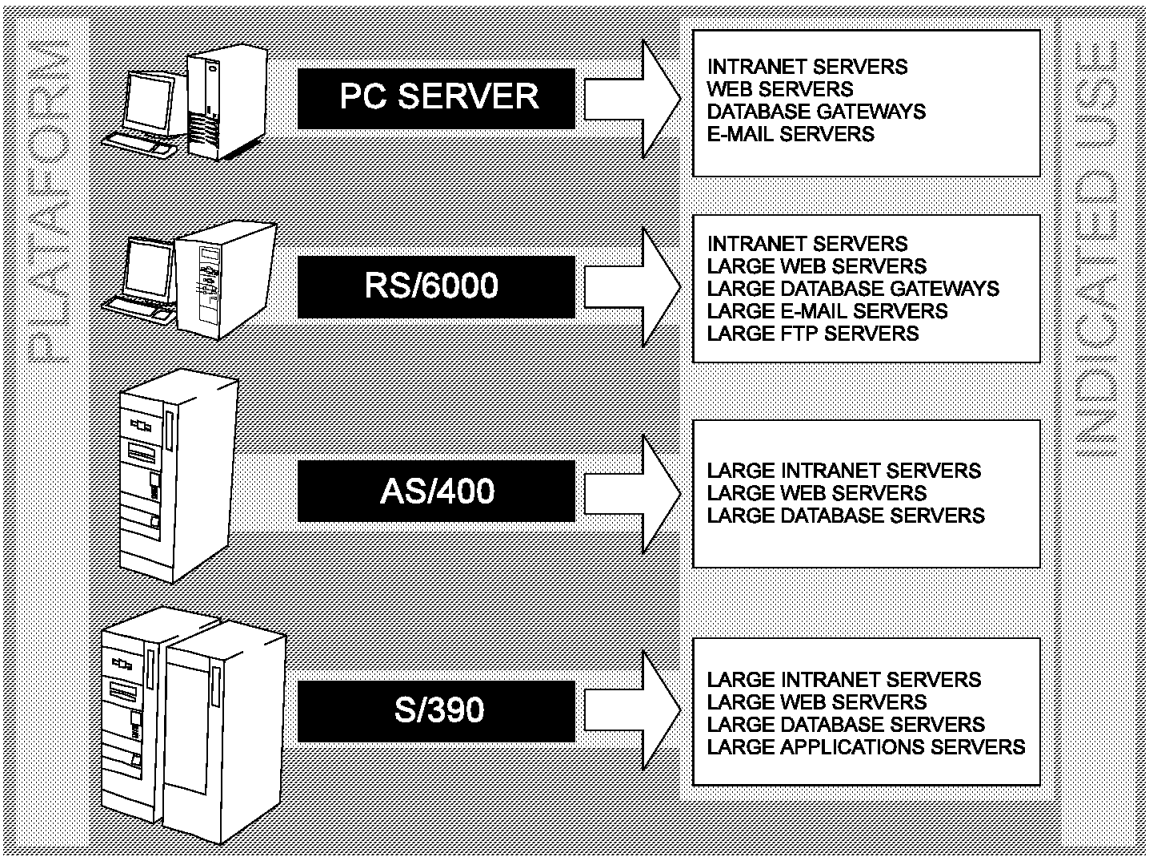


Figure 3. Platform and Service

1.3 Access Technologies

This area covers access technologies.

1.3.1 Spread Spectrum Technology

The wireless revolution will be driven by radio technology developed during World War II to protect military and diplomatic communications. From this cloak-and-dagger genealogy, spread spectrum radio is developing into a core technology for today′s wireless challenges. While available for many years, spread spectrum radio was employed almost exclusively for military use. In 1985, the FCC allowed spread spectrum′s unlicensed commercial use in three frequency bands: 902 to 928 MHz, 2.4000 to 2.4835 GHz and 5.725 to 5.850 GHz.

Spread spectrum radio differs from other commercial radio technologies because it spreads, rather than concentrates, its signal over a wide frequency range within its assigned bands. The two main signal-spreading techniques are direct sequencing and frequency-hopping. Direct sequencing continuously distributes the data signal across a broad portion of the frequency band. This technique modules a carrier by a digital code with a bit rate much higher than the information signal bandwidth. Frequency-hopping radios move a radio signal from frequency to frequency in a fraction of a second.

True to its military heritage, spread spectrum camouflages data by mixing the actual signal with a spreading code pattern. Code patterns shift the signal′s

frequency or phase, making it extremely difficult to intercept an entire message without knowing the specific code used. Transmitting and receiving radios must use the same spreading code, so only they can decode the true signal.

Obviously, spread spectrum radio is not the only wireless technology available. But in specific applications, its inherent attributes make it the technology of choice over traditional microwave radio or the optical technologies such as infrared and laser transmission, particularly in the last mile where wires can′t go or in hostile environment applications.

The most recent spread spectrum WAN/LAN developments have come through the integration of the radio with a full-function Ethernet bridge. A wide range of commercial spread spectrum products are being developed in response to the 1985 FCC Part 15 ruling. The key to commercializing spread spectrum is overcoming its complexity and cost. Most of the complexity in direct sequence radios resides in digital processing (DSP) or custom-designed chips. Today, all kinds of complex processing are available in the form of low-cost chips in everyday products. As practical commercial applications become better understood, spread spectrum will play an increasingly critical role in a world destined to depend on wireless technology.

There are some limitations when you use a spread spectrum link. You need to install the antennas on a configuration that must be on the same alignment, without any obstacle such as buildings, mountains, etc. If you have this kind of restriction, the solution is to install another set of antennas and radio modems to create a reflector node. This example is shown in Figure 4 on page 14 through Figure 8 on page 16.

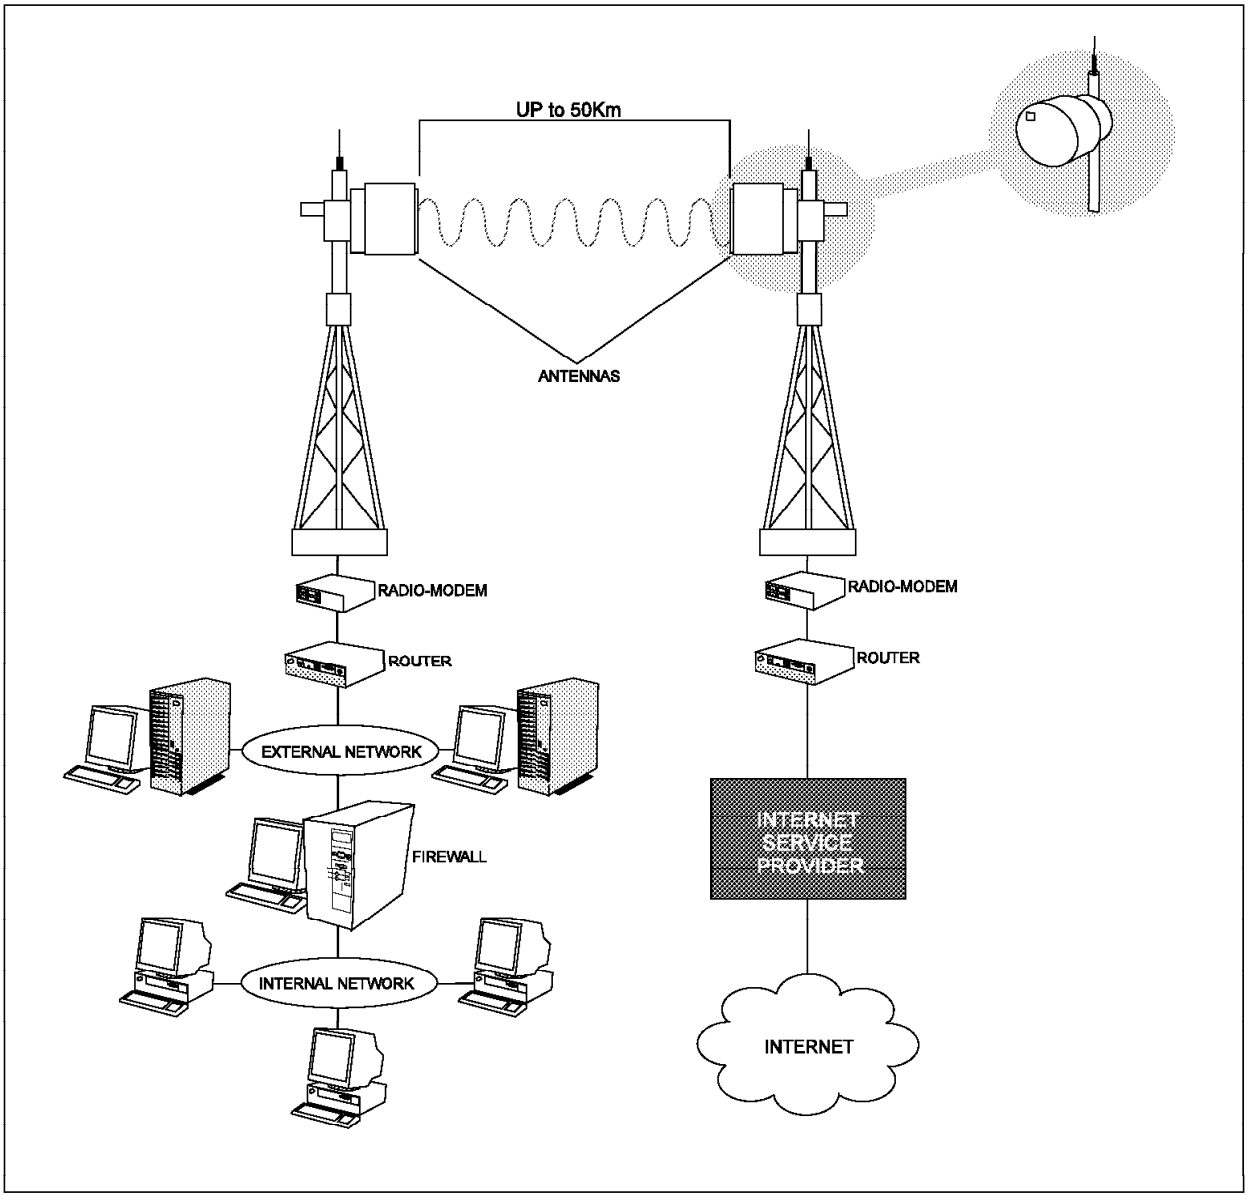


Figure 4. Spread Spectrum Link. This solution is excellent to connect networks to the Internet and to connect corporative sites. But there are some limitations such as the distance between the antennas and obstructions on the radio link flow, such as mountains and buildings. You can get high-speed connections, starting at 64 Kbps to 45 Mbps without spending money with a telecommunications provider services.

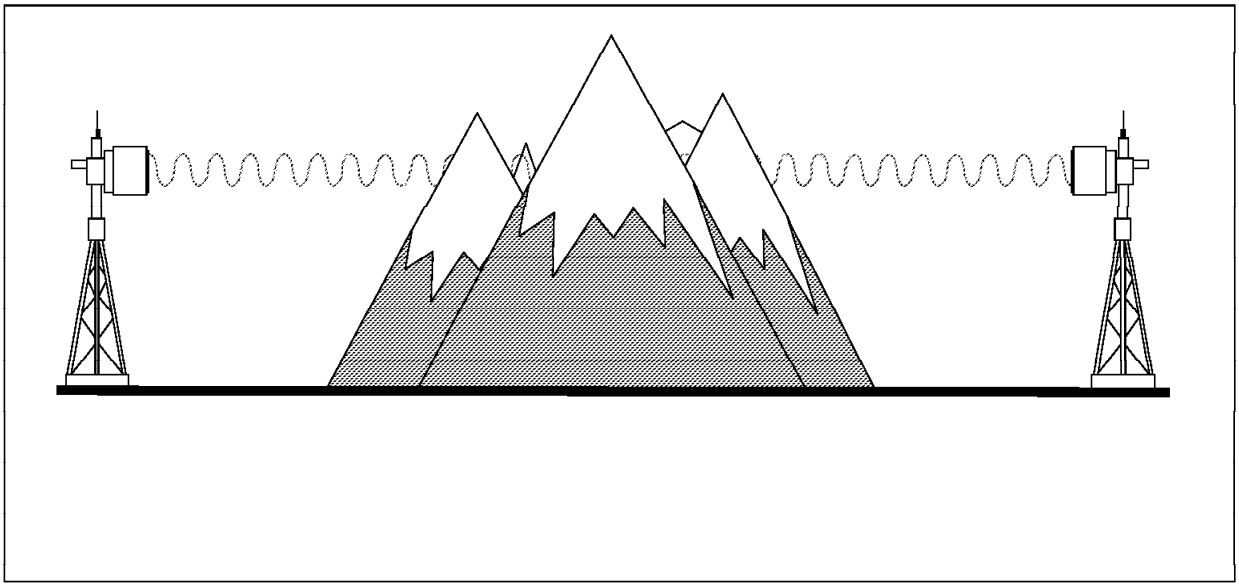


Figure 5. Natural Problems. Mountains and other natural obstructions are a problem to a spread spectrum solution.

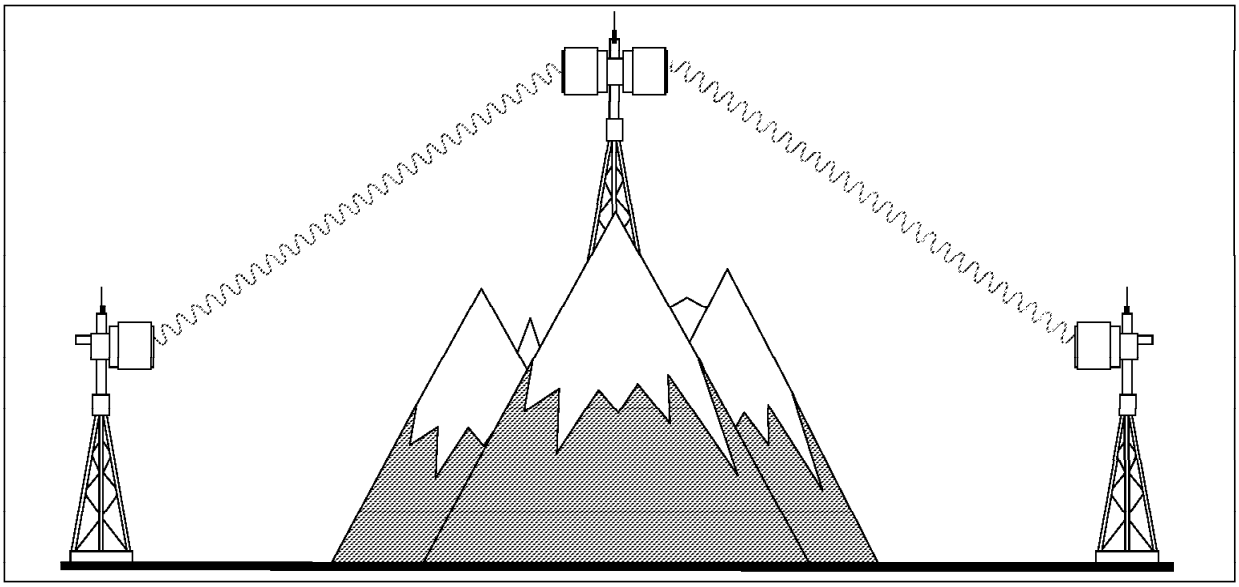


Figure 6. Solution. Using an additional set of antennas you can create a reflector to bypass the natural obstruction.

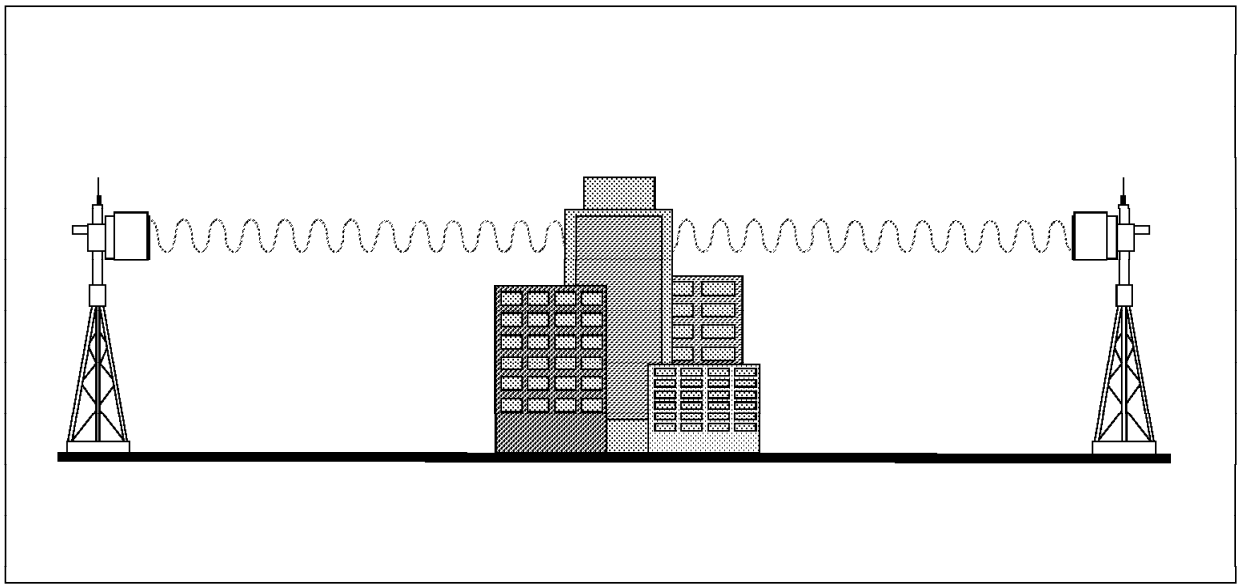


Figure 7. Big-City Problem. Large buildings, houses and towers can also interfere with spread spectrum transmission.

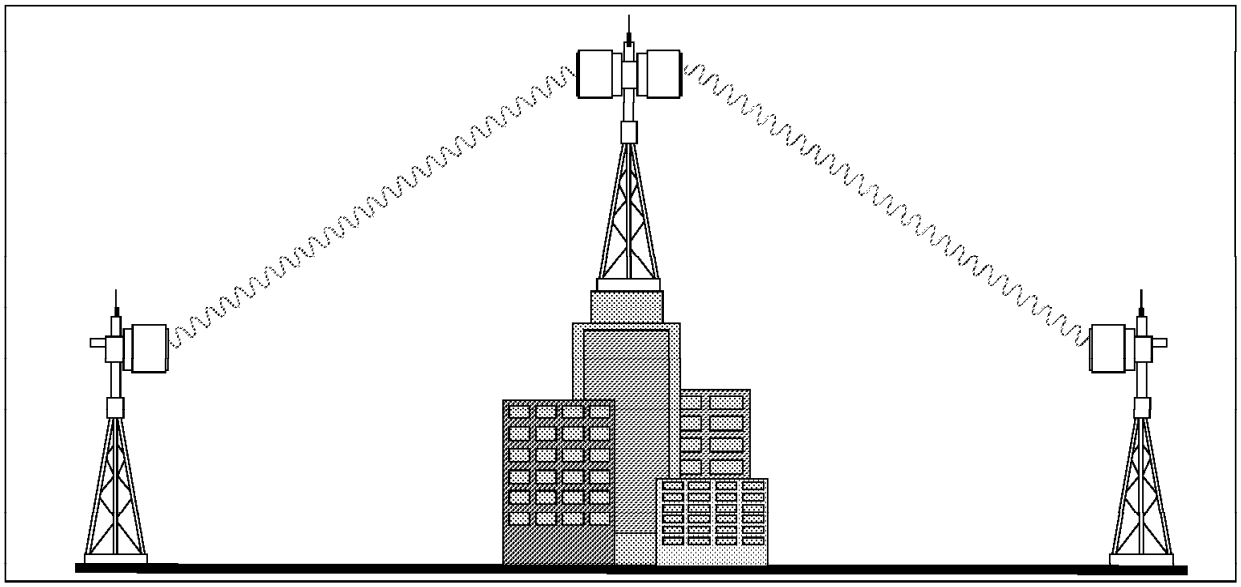


Figure 8. Solution. Like the natural obstruction solution you will need an additional set of antennas to create a reflector. Using a reflector you can bypass the obstruction and multiply the transmission range.

1.3.2 Leased-Line Connections

Leased lines are the most common way to connect a corporative environment to the Internet. They are stable and reliable. In some countries, you can get very cheap high-speed channels. There are many different kinds of leased connections. They can vary depending on the country, but the most popular speed and standards are as follows:

- 56 kbps: This is a digital phone-line connection capable of carrying 56,000 bps. At this speed, a megabyte will take about three minutes to transfer. This is 3.7 times as fast as a 14,400 bps modem.
- 64 kbps: This is also a digital phone-line connection capable of carrying 64,000 bps. At this speed, a megabyte will take about two minutes to transfer. This is 4.4 times as fast as a 14,400 bps modem.
- T1: This is a leased-line connection capable of carrying data at 1,544,000 bps. At maximum theoretical capacity, a T-1 line could move a megabyte in less than 10 seconds. That is still not fast enough for full-screen, full-motion video, for which you need at least 10,000,000 bps. T-1 is the most common speed used to connect networks to the Internet.
- T3: This is a leased-line connection capable of carrying data at 44,736,000 bps. This is more than enough to do full-screen, full-motion video.

1.3.3 Cable Modems

A *cable modem* is a device that allows high-speed data access (such as to the Internet) via a cable TV (CATV) network. A cable modem will typically have two connections, one to the cable wall outlet and the other to a computer (PC). Cable modem speeds vary widely. In the downstream direction (from the network to the computer), speeds can be anywhere up to 36 Mbps. Few computers will be capable of connecting at such high speeds, so a more realistic number is 3 to 10 Mbps. In the upstream direction (from computer to network), speeds can be up to 10 Mbps.

However, most modem producers will probably select a more optimum speed of between 200 kbps and 2 Mbps. In the first few years of cable modem deployment, an asymmetric setup will probably be more common than a symmetric setup. In an asymmetric scheme, the downstream channel has a much higher bandwidth allocation (faster data rate) than the upstream. One reason is that the current Internet applications tend to be asymmetric in nature. Activities such as World Wide Web (HTTP) navigating and newsgroups reading (NNTP) send much more data down to the computer than to the network. Mouse clicks (URL requests) and e-mail messages are not bandwidth-intensive in the upstream direction. Image files and streaming media (audio and video) are very bandwidth intensive in the downstream direction.

The fact that the word *modem* is used to describe this device can be a little misleading only in that it conjures up images of a typical telephone dial-up modem. Yes, it is a modem in the true sense of the word; it modulates and demodulates signals. But the similarity ends there because cable modems are practically an order of magnitude more complicated than their telephone counterparts. Cable modems can be part modem, part tuner, part encryption/decryption device, part bridge, part router, part NIC card, part SNMP agent, and part Ethernet hub.

Typically, a cable modem sends and receives data in two slightly different fashions. In the downstream direction, the digital data is modulated and then placed on a typical 6 MHz television carrier, somewhere between 42 MHz and 750 MHz. There are several modulation schemes, but the two most popular are QPSK (up to 10 Mbps) and QAM64 (up to 36 Mbps). This signal can be placed in a 6 MHz channel adjacent to TV signals on either side without disturbing the cable television video signals. The upstream channel is more tricky. Typically, in a two-way activated cable network, the upstream (also known as the reverse

path) is transmitted between 5 and 40 MHz. This tends to be a noisy environment, with lots of interference from HAM radio, CB radios and impulse noise from home appliances. Additionally, interference is easily introduced in the home, due to loose connectors or poor cabling. Since cable networks are tree and branch networks, all this noise gets added together as the signals travel upstream, combining and increasing. Due to this problem, most manufacturers will be using QPSK or a similar modulation scheme in the upstream direction, because QPSK is a more robust scheme than higher-order modulation techniques in a noisy environment. The drawback is that QPSK is slower than QAM.

There are several methods for computer connection, but it appears that Ethernet 10BaseT is emerging as the most predominant method. Although it probably would be cheaper to produce the cable modem as an internal card for the computer, this would require different printed circuit cards for different kinds of computers and, additionally, would make the demarcation between cable network and the subscriber′s computer too fuzzy.

The most popular service will undoubtedly be high-speed Internet access. This will enable the typical array of Internet services at speeds of 100 to 1000 times as fast as a telephone modem. Other services may include access to streaming audio and video servers, local content (community information and services), access to CD-ROM servers, and a wide variety of other service offerings. New service ideas are being born daily.

Cable modem pilot tests are already underway in many cable networks. But testing is still in an early phase, and large scale testing won′t take place until 1996. Many of the cable modems will first appear on the market in 1996. Wide scale deployments probably won′t start until some time in 1997.

There are many companies who are producing or have announced cable modem products. Included are: IBM, AT&T, COM21, General Instrument, HP, Hughes, Hybrid, 3COM, Intel, LANCity, MicroUnity, Motorola, Nortel, Panasonic, Scientific Atlanta, Terrayon, Toshiba, and Zenith.

As mentioned earlier, cable modems will enable data connections of much higher speeds than ISDN. ISDN transmits and receives at speeds of 64 kbps and 128 kbps. Cable modems will be able to receive data at up to 10 Mbps and send data at speeds up to 2 Mbps (some up to 10 Mbps). However, this is not the only advantage of a cable modem.

It is well known that the installation of an ISDN data connection for a residential subscriber is a very complicated process. The home user often has to act as his or her own system integrator. Installation requires careful integration of the telephone company service, the terminal adapter, the computer system, and the software. Service from the cable company will likely result in a technician bringing the modem to your home, installing the modem, installing the necessary software, and when the technician leaves your house, you will be up and operating. This places the installation and activation burden on the cable company rather than on the subscriber.

1.3.4 Integrated Services Digital Network (ISDN)

ISDN is an acronym for Integrated Services Digital Network. It is no longer necessary to use dedicated lines to gain the benefits of digital speeds or connectivity. The flow of digital information now begins at the user′s desktop and links it to the desktops of users around the world. From voice and data to complex images, full-color video and stereo quality sound, all are transmitted with digital speed and accuracy through what is now a totally digital network. ISDN replaces today′s slow modem technology with speeds of up to 128 kbps (kilobits per second) before compression. With compression, users in many applications today can achieve throughput speeds of from 256 kbps to more than 1,024 kbps, more than a megabit per second.

Digital lines are almost totally error free, which means that the slowdowns and errors typically encountered in today′s modern transmissions are no longer a problem. A single ISDN line can serve as many as eight devices: digital telephones, facsimiles, desktop computers, video units and much more.

Each device, in turn, can be assigned its own telephone number, so that incoming calls can be routed directly to the appropriate device. Any two of these devices can be in use at the same time for voice for data transmissions, and the lines can also be combined for higher data speeds. In addition, an almost unlimited number of lower-speed data transmissions (for e-mail, credit card authorization, etc.) can go on at the same time. In most cases, the same copper wires used today for what is typically called plain old telephone service can be used successfully for ISDN. This means most homes and offices are ISDN-ready today.

Often overlooked in the excitement of faster, more accurate data transmissions is the fact that ISDN represents the next generation of voice telephone service. It offers absolutely quiet, clear worldwide conversations every time plus a host of powerful call management and call handling capabilities. ISDN lines can be connected, interworked to virtually every other voice, data and packet network in the world, from a voice call across the street to a private terminal in a remote corner of the world; in short, ISDN lines are a faster, better, more economical way to communicate.

The 23B + D is an example of service configuration that provides 23 B channels and 1 D channel. The B channels carry user information such as voice calls, circuit-switched data, or video, while the D channel handles signaling information. When equipped, the D channel can control a maximum of 479 B channels. The B channel may be provisioned on the same facility as the D channel or on another Primary Rate Service T1 facility.

The basic Primary Rate Service (PRS) structure consists of 23 B channels and a D channel, for a total transmission rate of 1.544 Mbps, which is equivalent to a T1 facility. Each 64-kbps B channel carries user information such as voice calls, circuit-switched data, or video. The D channel is a 64 kbps channel that is used to carry the control or signaling information.

Single Line ISDN Service (SLS) is a platform-based switched digital service offering fast, flexible, highly reliable, and digitally clear connections with the simplicity of dialing a telephone. Based upon international communications standards, ISDN provides users access to the powerful capabilities of today′s Public Telephone Network for communicating across town or around the world. With Single Line ISDN Service, the same pair of wires that now delivers one

communication-at-a-time basic phone service to business or residence customers provides two primary, high-speed (64 kbps) communications channels that can be used simultaneously and independently to carry any combination of data, image, video, or voice calls. By combining these channels, data transfer at up to 128 kbps may be achieved. Single Line ISDN Service also provides a third auxiliary channel for low to moderate-speed data communications, which is ideal for point of sale, remote monitoring or telemetry applications. No special handling is required when voice calls are made between ISDN phones and conventional telephones; the network manages the necessary conversions. When conducting data calls, in order to utilize the B Channels for digital communications, ISDN-based equipment is required at both ends of the communications path, as is the case with conventional modem connections or fax machine transmissions. Certain ISDN equipment also allows for modem-to-modem communications.

Single Line ISDN Service includes a comprehensive 2B + D package. Contained in the standard package are numerous voice and data features. The standard features and functions support two terminals per basic rate service. Within the standard package there is limited flexibility for customization and various optional features can be added. Single Line ISDN Service does not offer B channel packet service capability.

The D or Delta channel carries signaling and/or packet data information, at speeds up to 16 kbps on basic rate service or Single Line ISDN Service, and signaling only information up to 64 kbps for primary rate service from the customer′s premises to the central office. The D channel has both data and signaling functionality; it does not have voice capability.

The B or Bearer Channel carries circuit-switched voice and/or data communications at speeds up to 64 kbps from the customers premises, over the loop facility, to the central office.

The B Channel circuit-switched data provides the capability of making data calls over the public switched network. Information is transmitted the same way as digitized voice. Like a voice call, a circuit-switched data call ties up network/system resources for the duration of the call. Similar to voice, calling line identification functionality is provided.

Chapter 2. Networking Hardware

This chapter presents the hardware commonly used in the Internet environment.

2.1 IBM 8235 Dial-In Access to LANs Server

The IBM 8235 Dial-In Access to LANs (DIALs) Server for token-ring and Ethernet is a dedicated multiport, multiprotocol remote access hardware server. This server supports remote personal computer (PC) users dialing into applications the same way users access applications from workstations directly attached to a token-ring or Ethernet local area network. With routing and bridging support for the following multiple protocols, a user can remotely access a variety of applications:

- NetBIOS for LAN servers
- IPX for NetWare
- 802.2 LLC for 3270 and SNA
- IP for TCP/IP applications
- AppleTalk Apple Remote Access (ARA) 2.0 (Ethernet Only)

Using standard dial networks, users with PCs and modems who are remote from the LAN can access LAN resources and work with applications as if they were working at locally attached LAN workstations.

Users in the field, such as agents, sales representatives, and employees who travel or work at home, have the ability to access their applications from any location that has dial-up telephone service. This extends the productivity of the workstation to the remote workplace. Using standard analog modems and dial-up telephone lines, the IBM 8235 and the IBM DIALs Client for OS/2, DOS, and Windows operating in the remote PC allow easy access to resources that users normally access from a workstation connected to a LAN. With support for multiple protocols and with high-performance filtering and compression techniques, excellent performance can be achieved when addressing a variety of applications remotely.

2.1.1 8235 System Components

The 8235 remote access system is made up of three basic components:

1. The Dial-in Access to LANs Client

This is a software application that runs on the remote PC providing the dial-in function. The DIALs Client supports DOS, Windows, and OS/2.

2. The 8235 Management Facility

This is a Windows application that allows the 8235 to be configured and managed from any LAN-attached workstation running IPX and Windows.

3. The 8235

This is a stand-alone hardware device that attaches to either a token-ring or Ethernet LAN and the public switched telephone network. The function of the 8235 hardware and its associated software is to:

- Provide physical attachment to the LAN and to eight modems

© Copyright IBM Corp. 1996

- Forward data from the LAN to the remote PCs and from the remote PCs to the LAN using any of the following protocols: IPX, IP, NetBEUI, AppleTalk ARA 2.0 and LLC
- Filter and compress data so as to minimize the amount of unnecessary traffic between the LAN and the remote PC
- Prevent unauthorized access to the LAN

2.1.2 Dial-In Access to LANs Server (DIALs) Client Software

DIALs Client is IBM′s multiprotocol dial-in software for workstations. It allows dial-in connections to any IBM 8235, providing full access to use any resources on a remote network. With the 8235 and its associated software (DIALs Client for OS/2, DOS, or WINDOWS), higher-level network applications treat the remote link as a local link. No custom applications are required to run remotely instead of locally.

Note

The DIALs Client is shipped with the 8235 with an unlimited right to copy.

DIALs Client contains the following software:

- OS/2 Drivers (NDIS and ODI)

These softwares provide support for OS/2-based communication programs. ODI can be provided with LAN adapter and protocol support (LAPS).

- DOS Drivers (NDIS and ODI)

These softwares provide support for DOS-based or Windows-based communication programs.

- Connect Application

This allows you to create, store and use connection files to dial in to remote networks from the OS/2, DOS, and Windows environments. The connect program:

– Provides traffic-flow statistics

– Displays error information

– Displays the modem status

– Displays the modem configuration

This section describes how to set up a connection to the Internet, via an IBM 8235 Dial-In Access to LANs Server, using the DIALs Client software for IBM OS/2 Warp Version 3.0 and OS/2 Warp Connect (DIALs Client/2 or DIALs Connect/2, both designations are correct). The DIALs Client software for Microsoft Windows 3.1 and 3.11 and Microsoft Windows for Workgroups 3.11 works essentially with the same windows and dialog boxes that the OS/2 version does. For any additional information about it, refer to *DIALs Client User′s Online Guide* in the IBM DIALs Program Group.

Figure 9 on page 23 shows the DIALs Connect/2 Version 4.02 product information.



Figure 9. DIALs Connect/2 Version 4.02 Product Information

The DIALs Connect/2 application manages the configuration of modems, phone numbers, passwords, and other items that establish the connection between the remote PC, the 8235, and the LAN. DIALs Connect/2 needs to be active only while connecting and disconnecting. However, it can remain loaded during the connection to provide information about the status of the call, traffic statistics, modem configuration, and more.

A separate connection file needs to be created for every access remote network users want to access. The connection file contains all of the Information DIALs Connect/2 needs to connect to the remote network. When a connection file for dialing in to a remote network is created, it should be saved and used each time the user wants to connect to that particular network. To run the DIALs Connect/2 application to create a connection file, the network administrator for the remote network must provide:

- The telephone number to dial
- A valid user name and, if required, a password
- The network protocols such as IPX, IP, and NetBEUI/LLC, that are required to make the connection

This section describes how to create a dial-in connection file, using the IP protocol, to access the Internet through a remote network.

To create and save a connection file:

1. Select **Connect/2** from the DIALs Connect/2 folder, as Figure 10 on page 24 shows.



Figure 10. DIALs/2 Folder

The DIALs Connect/2 window appears (see Figure 11).

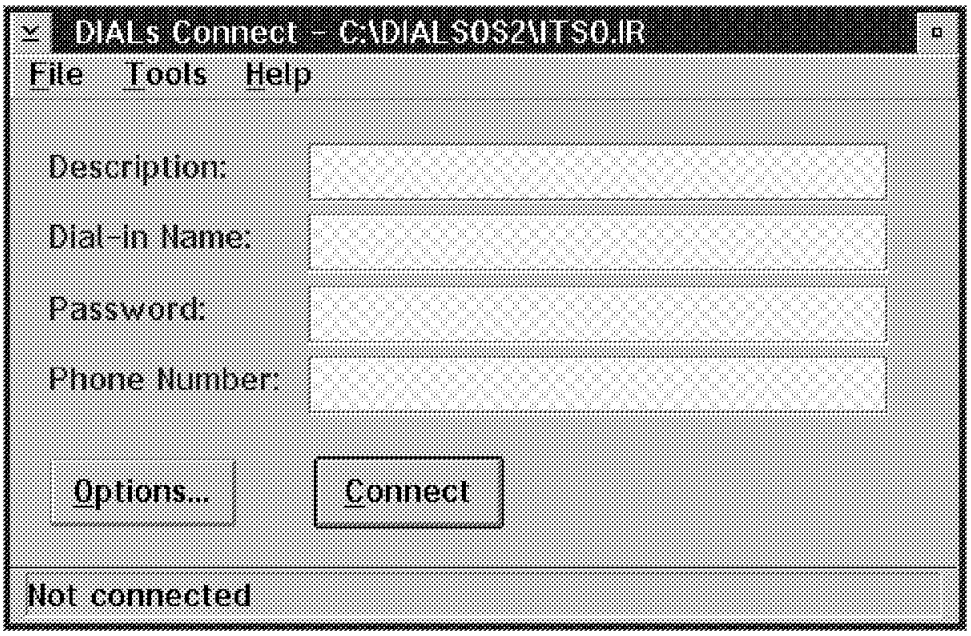


Figure 11. DIALs Connect/2 Window

Note

If the message DIAL.OS2 driver not loaded appears at the bottom of the DIALs Connect/2 window, make sure that the instructions to configure the drivers have been followed as described.

DIALs Connect/2 supports both NDIS (Network Driver Interface Specification) and ODI (Open Data-Link Interface) network protocol stack architectures. For each of these, DIALs Connect/2 contains a device driver (DIALNDIS.OS2 for NDIS, and DIALODI.OS2 for ODI) that provides the same software interfaces as LAN adapter device drivers to network program applications. Different OS/2 network applications require different network driver support, as illustrated in Table 5 on page 25.

Table 5. Common OS/2 Network Applications and Device Drivers

Network Application	Device Driver
LAN Services	NDIS
Communications Manager/2	NDIS
PC Support/2	NDIS
TCP/IP	NDIS
Novell NetWare	ODI
LAN Workplace	ODI

Although it is not possible to connect to a remote network unless the DIAL.OS2 driver is loaded, a connection file can still be created and saved.

2. Enter a description of this connection file in the Description box. This field is optional and can be up to 64 characters long (see Figure 12 on page 26).
3. Enter your dial-in user name provided by the network administrator in the Dial-in Name box. Dial-in user names are not case-sensitive and can be up to 64 characters long (see Figure 12 on page 26).

 Your dial-in user name is specific to the 8235 you are calling; it does not necessarily match your user name for using other services on the remote network such as file server or e-mail IDs.
4. If the network administrator has assigned you a password, enter it in the Password box. Passwords are not case-sensitive and are displayed as asterisks (*) when they are typed (see Figure 12 on page 26). Alternatively, enter the password when prompted for it during the connection process (see Figure 13 on page 26). For security reasons, passwords are not saved to the connection file.
5. Enter the telephone number of the remote network you are calling in the Phone Number box. Enter the number exactly as you would dial it manually, using up to 56 characters including commas and hyphens (see Figure 12 on page 26). Use commas if you need to add a pause (usually 2 seconds for each comma you use, but this varies with modem settings). Hyphens are optional. This allows you to enter long-distance prefixes and telephone company charge codes.

 Note: Do not include any modem dial commands, such as ATDT, in the Phone Number field.

Keep in mind that many modems cannot handle more than 36 characters for dialing, so that if DIALs Connect/2 reports an error while dialing, this might be the cause.

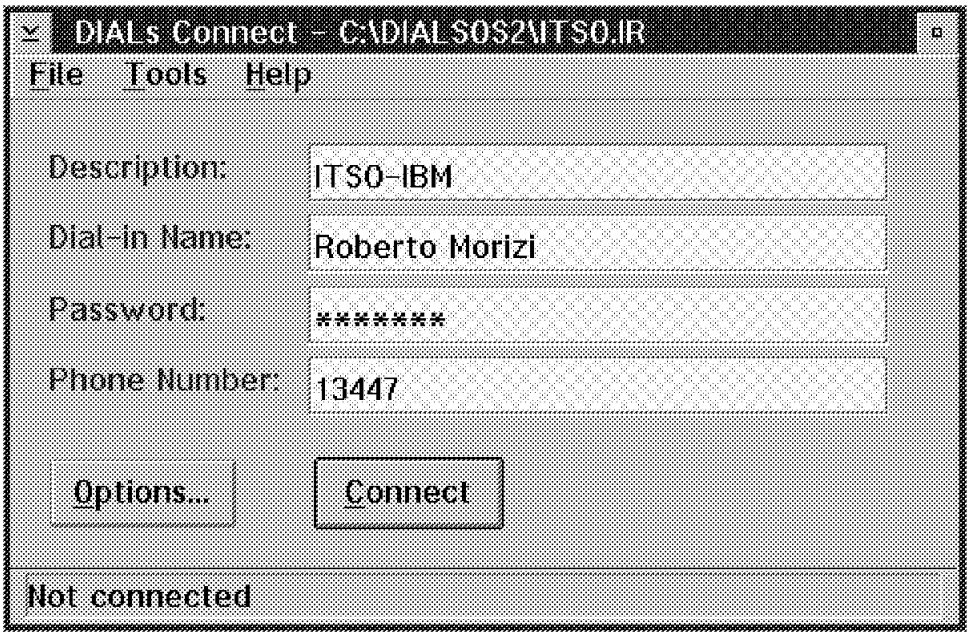


Figure 12. DIALs Connect/2 Window - Connection File

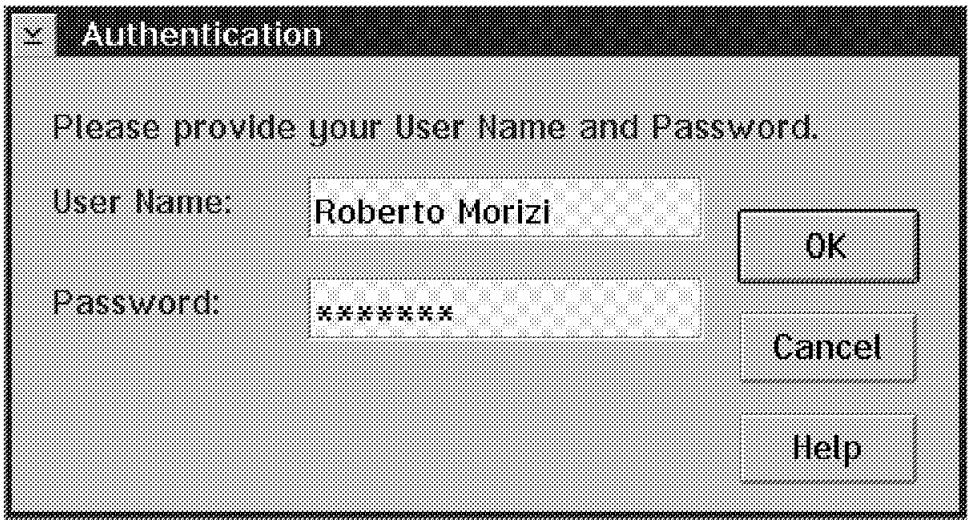


Figure 13. Authentication Window

6. Click on the **Options** button to set up the desired networking protocols and other features you want to use for this connection. The Connection File Options dialog box appears (see Figure 14 on page 27).



Figure 14. Connection File Options Dialog Box

7. Enable the network protocols you want to use when connected. It is possible to enable any combination of IPX, IP, NetBEUI, and LLC by selecting the check box next to each protocol. However, you will be able to use a selected protocol only if the remote server (8235) also supports that protocol. To disable a selected protocol, deselect its check box. To get access to the Internet, select **IP Protocol**.

 Table 6 lists common network applications and their corresponding protocols.

Table 6. Common OS/2 Network Applications and Protocols

Network Application	**Protocol**
LAN Services 3.0	NetBEUI/LLC
Communications Manager/2	NetBEUI/LLC
PC Support/2	NetBEUI/LLC
TCP/IP	IP
Novell NetWare	IPX
LAN Workplace	IPX

Note: When using the IP protocol, leave the IP Address field set to 0.0.0.0 unless the network administrator instructs you to enter an IP address. In

most cases, the dial-in workstation receives its IP address from the network, not from the value entered in this field.

8. If your user ID is set up on the 8235 to support roaming dial-back, select the **Request Roaming Dial-Back** check box.

 If this check box is selected, enter a phone number in the Dial-back Phone # field. Be sure that this is a valid telephone number for the telephone system used by the 8235. For example, if the 8235 must dial a 9 for an outside line, be sure to include that here. Roaming dial-back lets users tell the 8235 to call their modem back at a telephone number that they specify so they can reverse the charges for the telephone call. Not all 8235s support roaming dial-back, and not all users are set up to use this feature.

 For detailed information about IBM 8235's features, refer to *IBM 8235 Dial-In Access to LANs Server - Concepts and Experiences*, SG24-4816-00.

9. Select the **Connect automatically when connection file is loaded** check box to set up this connection automatically whenever this connection file is opened. If this option is not selected, you must click on the **Connect** button to make a connection after you open the connection file.

 Note: If you select this check box, you must make an icon for this connection file for DIALs Connect/2 to connect automatically. See *Creating an OS/2 Desktop Icon* in the DIALs/2 User's Guide. Figure 15 shows the DIALs/2 Folder and the new icon C:DIALSOS2ITSO.IR.



Figure 15. DIALs/2 Folder and the C:DIALSOS2ITSO.IR Icon

10. The Third-party security device installed selection tells DIALs Connect/2 to use a third-party security device that is set up on the 8235. If you select this check box, you will typically have to enter an additional password after connecting to the remote modem but before you have access to the 8235.

11. The Echo characters locally option tells DIALs Connect/2 to display characters on the screen as you type them. Select this check box only if you also selected the Third-party security device installed check box and the modem you are using does not echo keystrokes.

12. The Use default device option tells DIALs Connect/2 to use the default installed communications device or to override the device with another device.

13. Select **OK** to save the settings and return to the DIALs Connect/2 window. Select either **Save** or **Save as** from the File menu to save your configuration file (see Figure 16 on page 29).

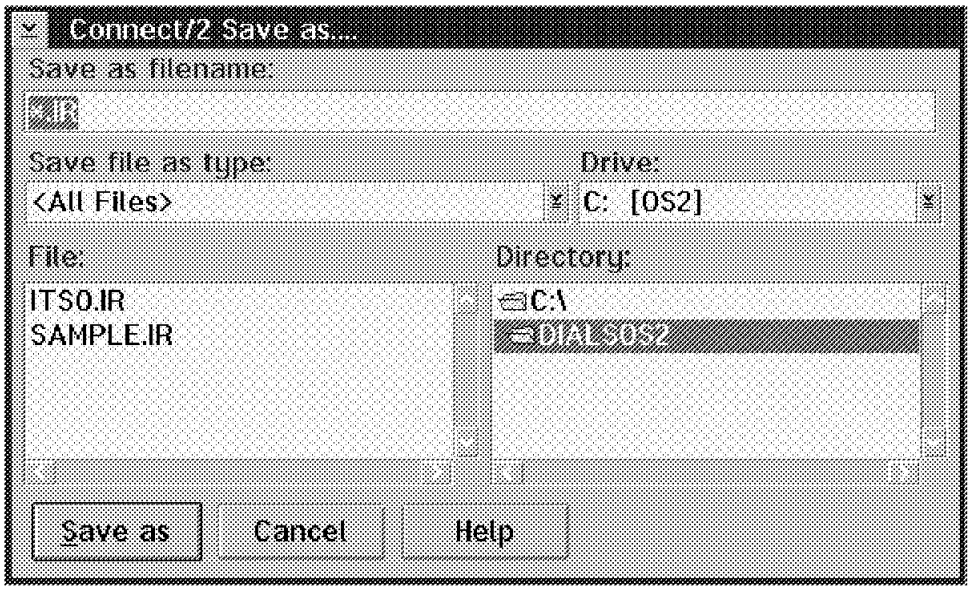


Figure 16. DIALs Connect/2 Save as Panel

14. The next step is to modify the port and modem settings. When you first install DIALs Connect/2, you need to set up the communications ports, telling DIALs Connect/2 what kind of modem or other communication devices you are using, as well as the COM port to which it is attached (or what drive to use in the event the communications device is not a COM port). You can also tell DIALs Connect/2 what speed to use for this connection (in bps), how to initialize the modem for the best possible connection, and so on. Use the Port Setup dialog box to modify all of these settings. Choose **Port Setup** from the Tools menu (see Figure 17).

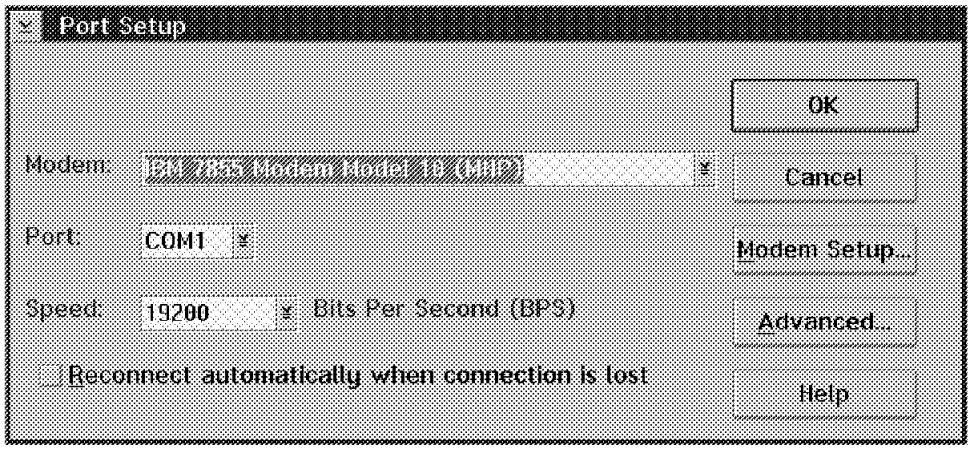


Figure 17. DIALs Connect/2 Port Setup Dialog Box

- Select the type of modem you are using from the Modem drop-down list.

 If the modem you want is not in the Modem drop-down list, click on **Modem Setup** to add your modem to the list (see Figure 18 on page 30).

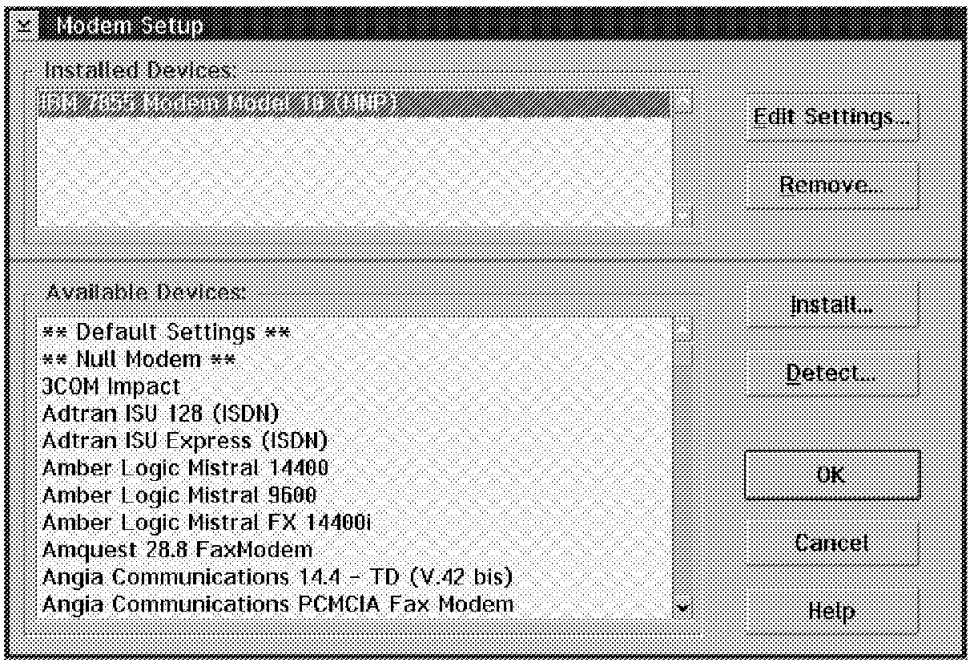


Figure 18. DIALs Connect/2 Modem Setup Dialog Box

If you need to set up a different communications device (modem or ISDN terminal adapter, for example), you can do so using the Modem Setup dialog box.

- To set up your communications device, select it from the Available Devices list.
- When the device you want is highlighted, click on **Install**. The device selected is added to the Installed Devices list.
- If you need to change the initialization string or other settings for your communications device from its default settings, select the device you added in the Installed Devices list and click on **Edit Settings** (see Figure 19).

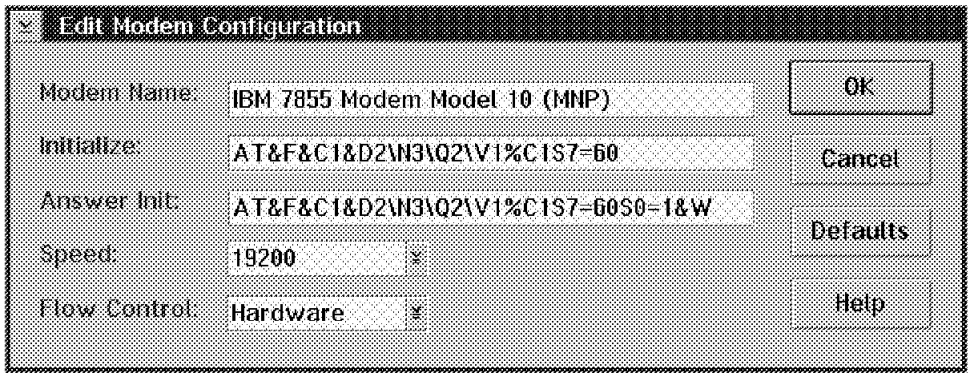


Figure 19. DIALs Connect/2 Modem Configuration Dialog Box

Use the Edit Modem Configuration dialog box to modify an existing modem configuration or create a new modem configuration.

- Modem Name Field: Lets you enter the name of the modem configuration you are currently adding or editing.

- Initialize Field: Contains the modem initialization string that DIALs Connect/2 sends to the modem to prepare it for a dial-in connection.
- Answer Init Field: Contains the modem initialization string that DIALs Connect/2 sends to the modem to prepare it to answer the telephone during a dial-back attempt.
- Speed Drop-Down List: Specifies the maximum speed at which a workstation can communicate with the modem in bps.

 To change the speed at which your workstation communicates with the modem, do not change this value; instead, change the value in the Speed field of the Port Setup dialog box.
- Flow Control Drop-Down List: Specifies the type of flow control the dial-in software uses (Hardware, Software, or None). Hardware flow control is also known as RTS/CTS. Software flow control is also known as XON/XOFF.
- Defaults Button: Restores the original configuration of the modem, discarding any changes that have been made. This button is active only if you have previously made changes to the configuration of a modem.

– Click on **OK** to close the Modem Setup dialog box and return to the Port Setup dialog box.

- Select the COM port to which the modem is attached from the Port drop-down list.
- Accept the default speed selected in the Speed drop-down list, or select another speed if you want.
- Select **Port Setup** from the System menu to verify the accuracy of your selections.
- If you want the DIALs Connect/2 software to automatically attempt to reestablish a lost modem connection, select the **Reconnect automatically when connection is lost** check box. Note that DIALs Connect/2 must be running at the time the connection was lost in order for the automatic reconnection to occur. If you do not select this check box, you are prompted to reconnect when the modem connection is lost.
- The default for settings are the most common ones. Click on **Advanced** to access the Advanced Settings dialog boxes. To change any of the default settings on the Advanced Port Settings dialog box, consult your system′s manual and the modem′s manual to verify your port settings.

 Note that there are two versions of the Advanced Settings dialog box: one if you are using a regular modem or ISDN terminal adapter, as Figure 20 on page 32 shows, and another if you are using the IBM WaveRunner digital modem, as shown in Figure 21 on page 33.

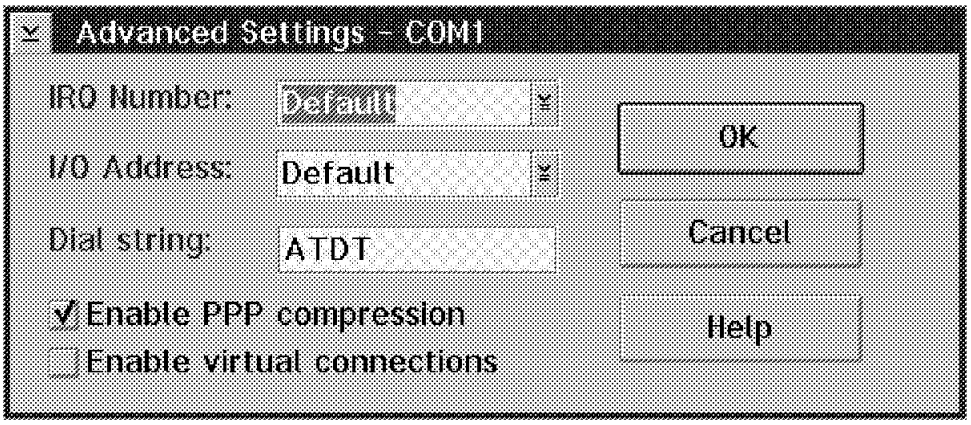


Figure 20. DIALs Connect/2 Advanced Port Settings Dialog Box

- IRQ Number Drop-Down List: If the COM port uses the standard IRQ number, leave this set to Default. If the COM port uses a non-standard IRQ number, use the drop-down list to select another value or enter that number here using a value between 2 and 15.
- I/O Address Drop-Down List: If the COM port uses the standard I/O address, leave this entry at Default. If the COM port uses a non-standard I/O address, use the drop-down list to select another value or enter that number here.
- Dial string field: In most cases, leave the values in the Dial String Field set to the default setting of ATDT. If the telephone connection requires pulse dialing, change the value to ATDP.
- Enable PPP Compression Check Box: This indicates whether DIALs Connect/2 and the 8235 should compress the information sent over the modem connection. This check box is selected by default. Also, if the 8235 has data compression enabled, selecting this check box can improve the speed of the dial-in connection. If the 8235 does not have data compression enabled, this setting is ignored.

 Note that DIALs Connect/2 must be dialing in to an 8235 with Version 3.5 or higher firmware installed for compression to be available.
- Enable Virtual Connections Check Box: This indicates whether DIALs Connect/2 and the 8235 should close your dial-in connection when you have not used the remote network for a certain length of time. This check box is not enabled in the default settings; you must enable the check box in order for virtual connections to be enabled.

 If this check box is selected and the 8235 (with Version 4.0 or higher firmware) has been configured to allow virtual connections, DIALs Connect/2 closes your dial-in connection when your workstation is idle (that is, when network access is not occurring) and re-open the connection automatically when network activity resumes.
- Click on **OK** to close the Advanced Port Settings dialog box and return to the Port Setup dialog box.

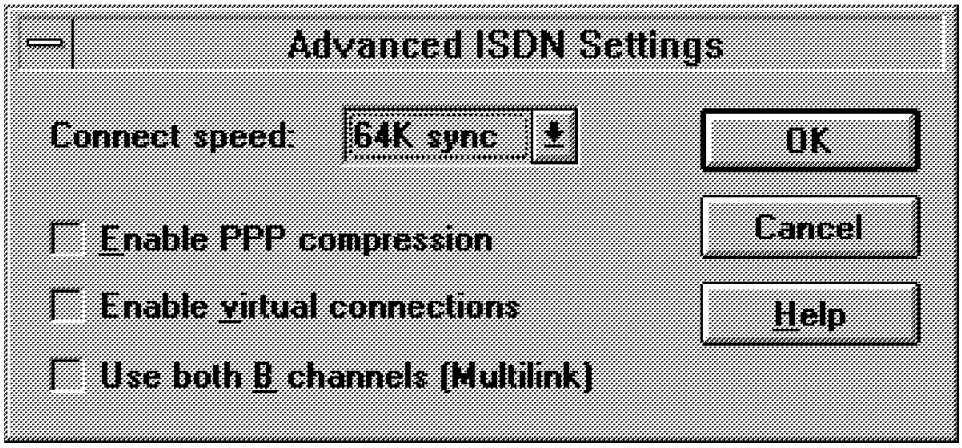


Figure 21. DIALs Connect/2 Advanced ISDN Settings Dialog Box

- Connect Speed: This indicates whether DIALs Connect/2 should attempt to connect at a speed of 64 kbps or 56 kbps. Your selection here depends on how your ISDN line was configured by your ISDN service provider.
- Enable PPP Compression Check Box: This indicates whether DIALs Connect/2 and the 8235 should compress the information sent over the connection. This check box is selected by default. If the 8235 has data compression enabled, selecting this check box can improve the speed of the dial-in connection. If the 8235 does not have data compression enabled, this setting is ignored.

 Note that DIALs Connect/2 must be dialing in to an 8235 running Version 4.0 or higher for compression to be available.
- Enable Virtual Connections Check Box: This indicates whether DIALs Connect/2 and the 8235 should suspend your dial-in connection whenever you have not used the remote network for a certain length of time, and resume it automatically when network activity resumes.
- Use Both B Channels (Multilink): This indicates whether DIALs Connect/2 and the 8235 should connect using MLP over your ISDN connection. This check box is not selected by default.

 If this check box is selected, you must be using the IBM WaveRunner digital modem to dial in to the remote network, and the 8235 on the remote network must also contain an 8235 BRI Module and have a working ISDN connection.

 Note

 DIALs Connect/2 provides support for high-performance channel aggregation using the industry-standard Multilink PPP Protocol (MLP). This feature allows dial-in connections to use multiple ISDN lines in a single connection session, providing increased bandwidth and performance.

- Click on **OK** to close the Advanced ISDN Settings dialog box and return to the Port Setup dialog box.

For additional information about IBM 8235 DIALs Client software for IBM OS/2 Warp Version 3.0 and OS/2 Warp Connect, refer to *DIALs/2 User's Guide* in the DIALs/2 folder.

2.1.3 Using the IBM Dial-Up for TCP/IP

IBM Dial-Up for TCP/IP allows you to use the Serial Line Internet Protocol (SLIP) or Point-to-Point Protocol (PPP) to connect to another TCP/IP host or to a service provider.

This section describes how to set up a connection to the Internet, via an IBM 8235 DIALs server, using the IBM Dial-Up for TCP/IP. We show a configuration using the Point-to-Point Protocol (PPP). For additional information, refer to *Introduction to TCP/IP* in OS/2 Warp's TCP/IP folder.

To configure dial-only connections for TCP/IP, installation of Multiprotocol Transport Services (MPTS) is required. See OS/2 documentation for information about installing MPTS.

To access the IBM Dial-Up for TCP/IP, select **Network Dialer** by double-clicking on its icon. Figure 22 shows the IBM Dial-Up for TCP/IP window.

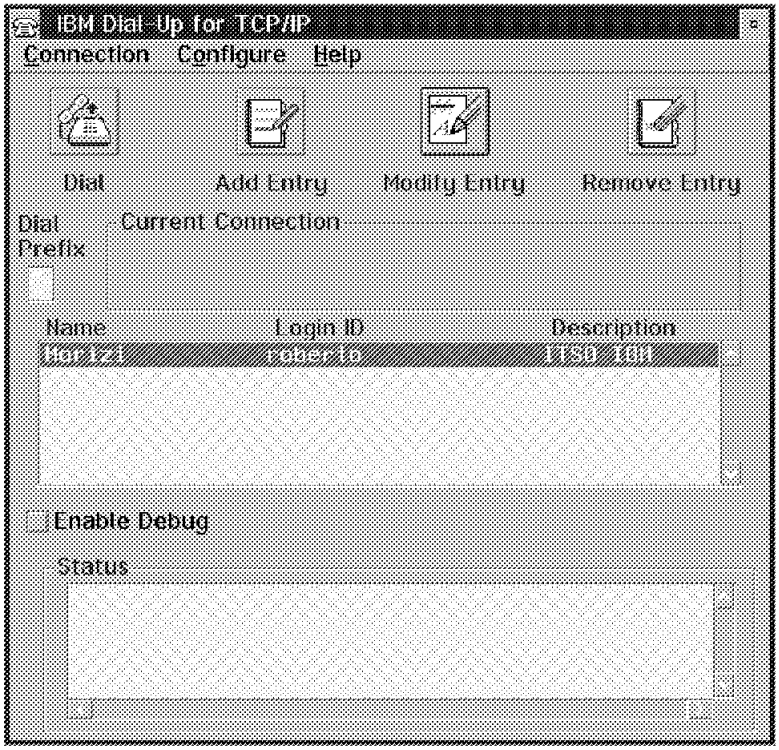


Figure 22. IBM Dial-Up for TCP/IP Window

- Dial/Hang-Up

 This push button changes depending on whether you have established a connection.

 Select **Dial** to establish the selected connection. Alternatively, you can select an entry and select **Dial** from the Connection pull-down menu. Select **Hang-Up** to close the connection. Alternatively, you can select **Hang-Up** from the Connection pull-down menu.

- Add Entry

 Select **Add Entry** to define a connection. Then, when the Add Entries window is displayed, enter the information to define the connection (see Figure 23).

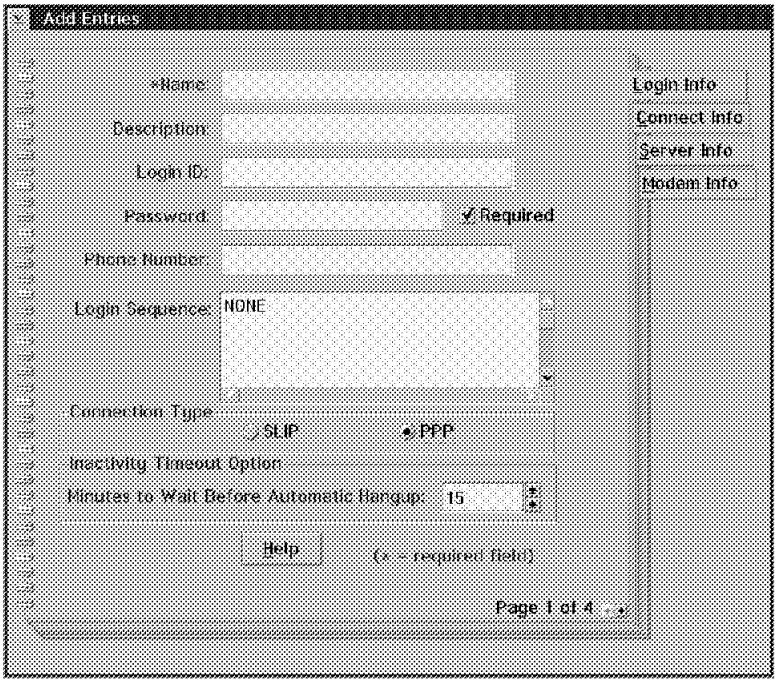


Figure 23. Add Entries Window

 - Name: Specify an identifier of the connection. This can be a comment or the name of a service provider. *This information is required*.
 - Description: Specify a description of the connection. Enter up to 11 characters.
 - Login ID: Specify the user identification assigned to you by the network administrator. This login ID is specific to the 8235 you are calling; it does not necessarily match your user name for using other services on the remote network such as file server or e-mail IDs. Logon IDs are not case-sensitive.
 - Password: Specify the password assigned to you. Passwords are not case-sensitive and are displayed as asterisks(*) when they are typed.
 - Phone Number: Specify the phone number used to access the destination host or service provider′s network; include any long-distance access codes and the area code.

 Note: Do not include any modem dial commands, such as ATDT, in the Phone Number field.
 - Login Sequence: Specify the login sequence that you want to use, if any. You can use a login sequence to automate a connection.

 To accommodate a variety of connection sequences, this field may contain:

 - The reserved word NONE. This indicates no login sequence is required beyond the physical modem connection.

- Blank, or no entry. If this field is left blank, and the Login ID and Password fields are filled in, then when IBM Dial-Up for TCP/IP receives the login sequence:

 login:

 password:

 The contents of the Login ID and Password fields are sent in response.

- The name of an ASCII or REXX connection script (or response file). This file is executed at connection time to negotiate the modem setup, dial to the destination host, and log into the host.
- A login sequence, which consists of a series of send-expect verbs.

Information entered in this field is stored in the TCPOS2.INI file.

If you are using a service provider, each provider may use a slightly different sequence for establishing a connection. You must tailor your login sequence to match each service provider.

– Connection Type: Select either SLIP or PPP if you are using the Serial Line Internet Protocol (SLIP) or Point-to-Point Protocol (PPP) to connect to the IBM 8235 DIALs Server.

– Inactivity Timeout Option: Specify the amount of idle time (in minutes) to be allowed before IBM Dial-Up for TCP/IP closes the connection.

• Modify Entry

Once you have defined a connection, select **Modify Entry** to change the definition of a selected connection (see Figure 9 on page 23 6.).

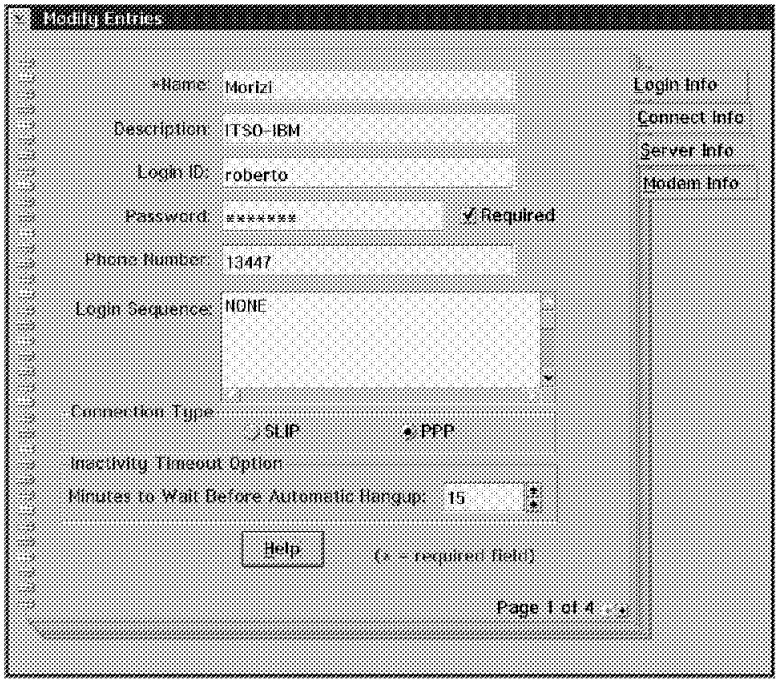


Figure 24. Modify Entries Window / Login Info Window

This first Modify Entries window shows the login information.

The Connect Info window allows you to configure the following information (see Figure 25 on page 37):

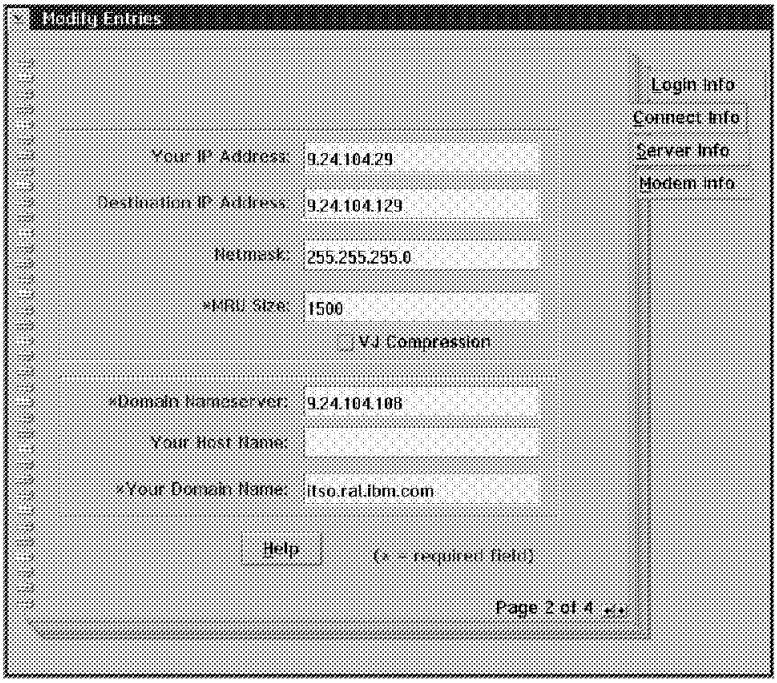


Figure 25. Modify Entries Window / Connect Info Window

- Your IP Address: Specify the 32-bit dotted decimal notation Internet Protocol (IP) address assigned to you.
- Destination IP Address: Specify the 32-bit dotted decimal notation Internet Protocol (IP) address of the destination host to which you want to connect (such as the IBM 8235 DIALs Server's IP address).
- Netmask: Specify the 32-bit dotted decimal notation network mask (subnet) used to indicate which portion of your IP address represents the network address and which represents the host address.
- MTU or MRU Size: Specify the MTU or MRU that your connection can handle. This is the largest possible unit of data that can be sent on a given medium in a single frame. If you are using SLIP, the default is 1006. If you are using PPP, the default is 1500. Valid values range up to 1500. *This is a required field.*
 - MTU - Maximum Transmission Unit
 - MRU - Maximum Response Unit
- Domain Name Server: Specify the 32-bit dotted decimal notation Internet Protocol (IP) address of the server that resolves host names to IP addresses. *This is a required field.*
- Your Host Name: Specify the symbolic name assigned to your computer.
- Your Domain Name: Specify the name of the domain in which your computer resides. The domain name includes all subdomains and the root domain separated by periods. *This is a required field.*

After you have entered the information on this page, select the **Server Info** tab.

The Server Info window allows you to configure the following information (see Figure 26 on page 38):

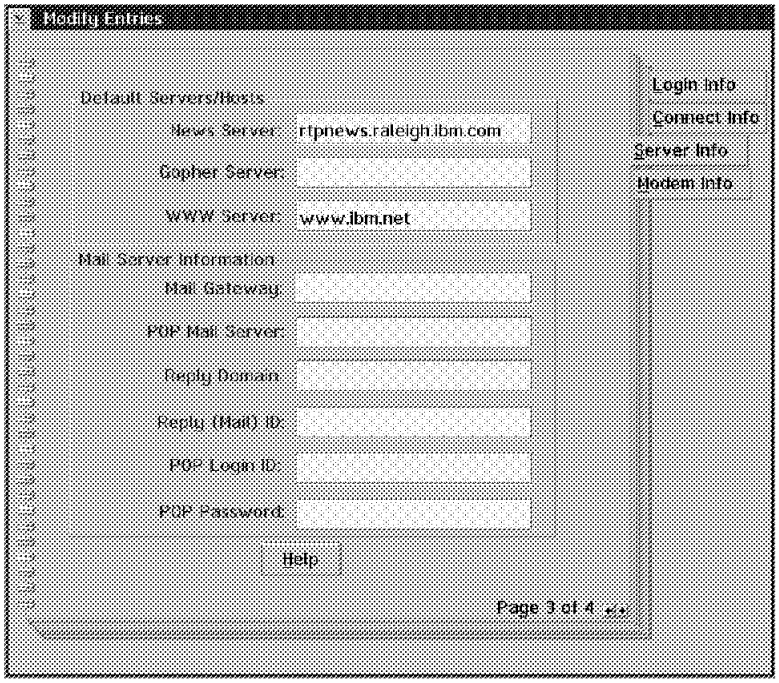


Figure 26. Modify Entries Window / Server Info Window

– Default Servers/Hosts

 - News Server: Specify the host name or IP address of the default news server.
 - Gopher Server: Specify the host name or IP address of the default Gopher server.
 - WWW Server: Specify the host name or IP address of the default World Wide Web (WWW) server.

– Mail Server Information

 - Mail Gateway: The mail gateway routes the mail to the recipients. The mail gateway is analogous to a POP server. By default, the entry for the POP mail server field is used as the entry for the mail gateway field. The mail gateway field cannot use an IP address, so it is recommended that you use a host name for the POP mail server field.
 - POP Mail Server: Specify the host name of the default mail server.
 - Reply Domain: Specify the name of the domain in which your mail server resides. The domain name includes all subdomains and the root domain separated by periods.
 - Reply (Mail) ID: Specify the identifier assigned to you for use in sending and receiving e-mail.
 - POP Login ID: Specify the identifier assigned to you for access to the mail server.
 - POP Password: Specify the password assigned to you for the mail server.

After you have entered the information on this page, select the **Modem Info** tab.

The Modem Info window allows you to configure the following information (see Figure 27):

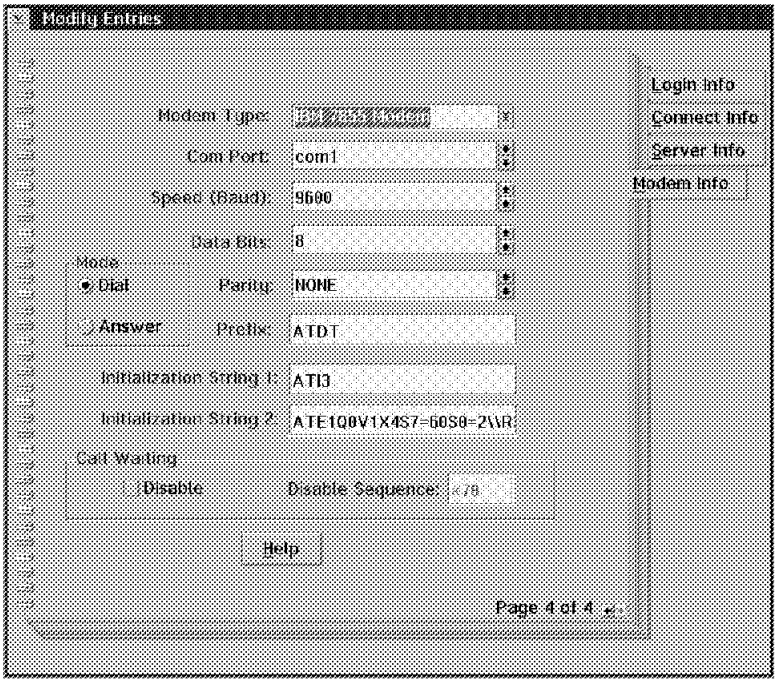


Figure 27. Modify Entries Window / Modem Info Window

- Modem Type: Specify the type of modem you are using.
- COM Port: Specify the name of the communications port of your computer to which your modem is attached. The default communications port is COM1.
- Speed (Baud): Specify the speed of the connection. This may be equal to or less than the capacity of your modem. The default speed is 9600 bps. Valid speeds are from 1200 to 115200 bps (async-to-modem bit rate).
- Data Bits: Specify the number of data bits in each character sent or received. Valid values are 7 and 8. The default is 8.
- Parity: Specify the parity of the connection. A parity bit is appended to a group of binary digits to cause the sum of the digits to be either even or odd. This parity bit is used in parity checks and should match the setting of the receiving modem.

 Valid values are NONE, SPACE, MARK, EVEN, and ODD. The default is NONE.
- Prefix: Specify the dial prefix for your modem. This is the attention command string that is passed to the modem and that preceeds the phone number. The default in Dial mode is ATDT. The default in Answer mode is ATS0=2S7=30. This information should be supplied in your modem documentation.
- Initialization String 1: Specify the initialization string for your modem. This is the command that initiates the modem. This information should be supplied in your modem documentation.

- Initialization String 2: Specify the initialization string for your modem. This is the command that initiates the modem. This information should be supplied in your modem documentation.
- Call-Waiting: If your phone service includes call-waiting, you will want to disable call-waiting while you are using the modem. If you disable call-waiting, you must also specify a Disable Sequence.

 If you have chosen to disable call-waiting, specify the phone key sequence used to disable this service. This information is required to disable call-waiting and can be found in your phone book.

To save your connection information, select the **Save** push button in the Closing Dial Configuration window. If there are required fields that are not complete, an Entry Input Error message appears and you are taken to the field that has the error.

- Remove Entry

 Select **Remove Entry** to delete the definition of the selected connection. The definition is deleted and the entry is removed from the connection list. Alternatively, you can select **Remove Entry** from the Configure pull-down menu.

To establish a connection, select an entry from the connection list and select the **Dial** push button on the IBM Dial-Up for TCP/IP window. Alternatively, you can select an entry and select **Dial** from the Connection pull-down menu. Figure 28 shows the information you will receive after establishing the connection.

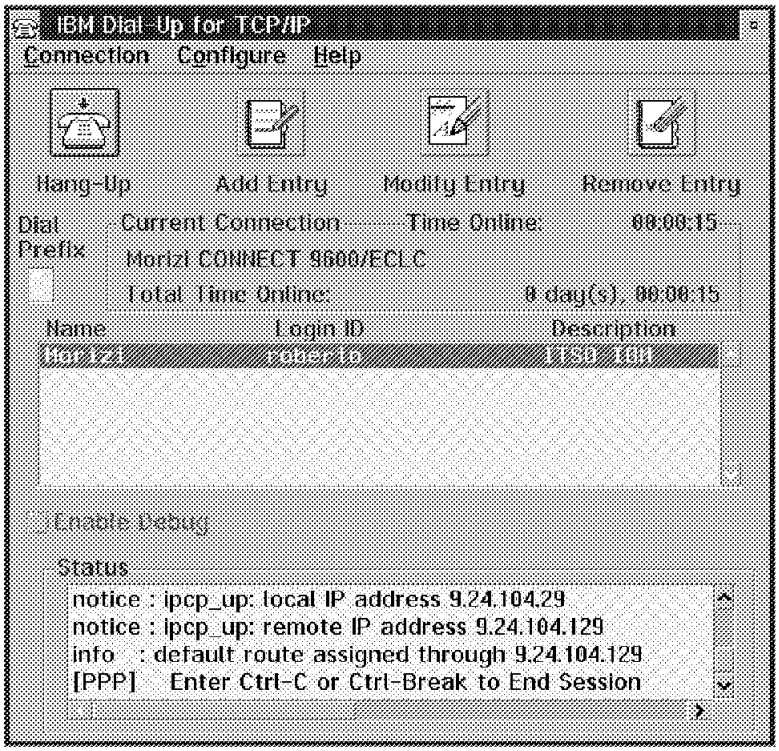


Figure 28. IBM Dial-Up for TCP/IP / Connection Status

Note

If your workstation has both local and remote access and, after dialing and connecting to the IBM 8235 DIALs server, you cannot load the IBM WebExplorer and even ping the 8235, the name server or any of your LAN′s routers through an OS/2 window, take a look at your workstation′s routing table to check if the IP addressing is compatible to the access type you are using, local or remote.

Then you can, for instance, access the Internet World Wide Web using the IBM WebExplorer.

2.1.4 IBM 8235 New Features

This section describes the new features provided by DIALS Release 2.0 and DIALS Release 4.0.

2.1.4.1 DIALS Release 2.0

1. Dial-In:

For the dial-in function, 8235 Version 2.0 provides the following features:

- ARA 2.0 dial-in support for Ethernet 8235s. (ARA 1.0 dial-in is not supported.) ARA dial-in provides the following features:
 - IP forwarding (MacTCP)
 - Routing or end-node forwarding support for ARA clients
 - AppleTalk device and zone filtering per user, per port, or per 8235
- Simultaneous PC dial-in over Point-to-Point Protocol (PPP) for the following protocols:
 - NetWare Internet Packet Exchange (IPX support)
 - Transmission Control Protocol/Internet Protocol (TCP/IP)
 - NetBIOS Extended User Interface (NetBEUI)
 - 802.2/Logical Link Control (LLC) (SNA)
- Support for the Novell Client for DOS/Windows, or Virtual Loadable Modules (VLMs)
- Windows for Workgroups (WFW) 3.11 support

2. Shared Dial-Out Access

 This is used for access to external asynchronous services such as CompuServe.

3. LAN-to-LAN Support
 - Connections between two networks routing any combination of TCP/IP and IPX over a dial-up link. AppleTalk LAN-to-LAN routing is supported for the Ethernet models of the 8235.
 - Connection features including idle detect, persistence, back-up telephone numbers, dial back, and timed connections.
 - LAN-to-LAN connections established automatically or via the command shell (scripting possible).
 - Leased-line support.

- AppleTalk device and zone filtering for the Ethernet models of the 8235.

4. Centralized Management
 - All protocols and features are manageable from the 8235 Management Facility for Windows.
 - Management Facility tuning for large IPX networks.
 - BOOTP/TFTP automatic downloading.
 - Command shell via IP Telnet, or dial-in on a PC.
5. Additional Security
 - Security Dynamics ACE/Server (SecurID) support for multiprotocol dial-in.
 - NetWare Bindery authentication for all protocols, including ARA 2.0.
 - 8235 user list.
 - Roaming or fixed dial back.

Note

Release 1.1 and 1.0 DIALs Client for OS/2, DOS, and Windows software is compatible with all 8235 models and previous releases, including Release 2.0. The new DIALs Client software Release 2.0 is shipped with 8235 Release 2.0 and is available in an upgrade kit for previous 8235 models. DIALs Release 2.0 Client software is not compatible with previous models of the 8235, unless the models are upgraded to microcode Release 2.0.

2.1.4.2 DIALS Release 4.0

1. Dial-In
 - Multiprotocol Support: Simultaneous multiprotocol dial-in over PPP: IPX (VLMs and NETX supported) TCP/IP, NetBEUI, 802.2/LLC.
 - VxD Windows Client Feature Summary: Client has been re-designed to enable support for:
 - Windows Virtual Device Driver VxD that only uses 2 KB of client conventional DOS memory (versus 34 KB)
 - Multilink PPP protocol (MLP)
 - Channel aggregation (2B)
 - Stac 4.0 compression
 - Port driver for internal ISDN adapters (digital modems, TAs)
 - Native driver support for IBM WaveRunner digital modem
 - New port driver programming interface (API)
 - Virtual connections
 - New intelligent setup facility
 - Easy Client installation scripting
 - Client event logging application
 - Virtual Connections: The ability to automatically suspend and resume a physical connection while spoofing network protocols, routing and applications. The physical connection is only brought up on demand.

- Spoofing: When a virtual connection is suspended, the ability for a device to determine what is not meaningful traffic. Rather than establishing the connection, the device responds to the source of the traffic with the response that would have been generated by the intended destination device.
- Dial-in Channel Aggregation: The ability to use more than one communications channel per connection. By aggregating both 64-kbps ISDN B-channels users can take advantage of 128-kbps dial-in connections. Fast 128-kbps data transfer rates reduce large file transfer times.
- IBM WaveRunner Digital Modem (Internal ISDN terminal adapter): Provides support for the ISA and PCMCIA versions of the IBM WaveRunner digital modem. The three supported modes are Async V.32 bis modem, ISDN V.120, and Sync Clear Channel.
- Easy client Setup:
 - An intelligent client setup program that includes a Connection File Wizard that walks the user through the installation and modifications to client software.
 - The ability to automatically detect attached communications adapters.
 - Powerful file copy mastering capability.
 - Client event logging application provides extensive troubleshooting information. Log information can be displayed to the screen or to a file.
- Power Switching: Allows users to switch back and forth between communications adapters. This is perfect for employees who use one type of communications adapter when working at home (ISDN) and another adapter (V.34 modem) when traveling.
- Express Installation: A new client installation scripting that enables network managers to establish defined defaults that make client installation and deployment easier.
- Third-Party Client Support: Dial-in access from Windows 95 and Windows NT 3.5, Apple′s ARA, and IBM′s OS/2 DIALS.

 Customers using Windows 95, Windows NT, MAC OS or OS/2 can seamlessly use an IBM 8235 as their dial-in server.
- Client Event Logging Application: Events can be displayed on the screen and/or saved in a text file. The logged events include:
 - Buffer allocation/management
 - PPP events and state transitions
 - PPP negotiation options
 - All frames transmitted and received
 - Multilink (MLP)
 - Compression
 - Network protocol decoding (basic IPX, IP and NetBEUI frames)
- New Port Driver: The new port driver provides support for internal client ISDN terminal adapters such as the IBM WaveRunner.

Internal ISDN adapters eliminate the async-to-sync conversion overhead required by external terminal adapters.

2. New Application Programming Interface (API): The IBM DIALs 4.0 port driver API enables third parties to independently develop IBM DIALs drivers for their hardware. Many internal ISDN terminal adapters do not present a standard PC 8250/16450/16550 UART interface.
3. Enhanced Stac 4.0 Compression: IBM upgraded the Stac compression algorithm from 3.0 to 4.0. Stac 4.0 is faster and more memory efficient. For digital terminal adapters where there is no compression done by the ISDN TA or X.25 PAD, it is essential that the compression algorithm used on the client be as lean and fast as possible.
4. LAN-to-LAN Features
 - Virtual Connections (VC): The ability to automatically suspend and resume a physical connection while spoofing network protocols, routing and applications. The physical connection is only brought up on demand.
 - Spoofing: When a virtual connection is suspended the ability for a device to determine what is not meaningful traffic. Rather than establishing the connection, the device responds to the source of the traffic with the response that would have been generated by the intended destination device. Spoofing is done for file server connections (NetWare drive mapping), routing tables (IP RIP and IPX RIP), SAP tables, TCP connections, and SPX connections.
 - Floating Virtual Connections (FVC): The ability to resume a suspended virtual connection on a port other than the port on which the original virtual connection was established. It can reduce the need to dedicate ports to specific users.
 - Juggling Virtual Connections (JVC): The ability to have more suspended virtual connections than there are ports on the IBM 8235. Customers can have many more suspended users than they have ports. JVC maximizes the utilization of server communications ports.
 - Persistent Connections (PC): An IBM 8235 configuration option that allows the server to re-establish the connection in the event of an unexpected line drop.
 - Timed LAN-to-LAN Connections (TLC): The ability for network managers to schedule LAN-to-LAN connections (for example, establish a LAN-to-LAN connection at 10 am and terminate the connection at 1 pm).
 - Piggybacking Updates: A virtual connection synchronizing mechanism where routing update messages are sent across the link only when the link is open for real data traffic.
 - Timed Updates: A virtual connection synchronizing mechanism where at a specified interval the suspended virtual connection is resumed to enable routing update messages to be sent across the link.
 - Triggered Updates:
 - A virtual connection synchronizing mechanism where routing update messages are sent across the link only when there is a RIP or SAP database change.
 - Triggered update setup options include additions only, deletions only, or additions and deletions.

- Channel Aggregation (Multilink PPP, MLP): The ability to use more than one communication channel per connection. LAN-to-LAN connections can aggregate all IBM 8235 channels (analog or digital) up to the number of ports on the server.
- Packet Fragmentation: The ability to configure a default packet size over which packets will be fragmented for more efficient distribution over aggregated communications links.
- LANConnect Applets: LANConnect applets for both PC and MAC allow for scripting of on-demand LAN-to-LAN connections.
- Delta Technology: Specialized remote adaptive routing protocols for optimizing bandwidth. It prevents unnecessary traffic from being sent over slow WAN connections by only sending the changes (deltas).

5. Management and Security Features
 - PC and MAC Server Management: Protocols and features can be managed by MAC or Windows versions of IBM NetManager (MAC Appletalk, PC/Windows IPX and IP).
 - IP Download: IBM MF will be able to download new code images and configurations when running over either IP or IPX protocol stack.
 - SNMP Management: MIB II and others.
 - Security: Provides support for agent software from Security Dynamics and Digital Pathways. Centralized authentication via IBM user list, NetWare Bindery, TACACS and most third-party hardware security solutions are supported.

2.1.5 What Is a Virtual Connection?

A virtual connection is a standard LAN-to-LAN or PC single-user dial-in connection that is enhanced to detect when no meaningful traffic has been sent over the connection for a period of time, at which time the physical connection is suspended while network protocols (IPX and TCP/IP) are spoofed by devices at either end of the connection. Subsequently, when meaningful traffic is received by either of the devices, the physical connection is automatically resumed and the data is forwarded over the communications link. Virtual connections minimize connect-time costs by physically disconnecting the circuit when there is no meaningful traffic.

Another benefit of a virtual connection is ease-of-use and management. Once the original connection is established, no user or system administrator intervention is required. The physical link is automatically suspended and resumed on demand.

2.1.6 What Is Channel Aggregation?

New high-performance channel aggregation technology enables dial-in and LAN-to-LAN users to establish more than one communications channel per connection. IBM channel aggregation technology utilizes the industry-standard protocol known as Multilink PPP for maximum client/server device interoperability and investment protection. Packet fragmentation is also available for maximum performance.

2.1.7 8235 Management Facility

The 8235 Management Facility is a device management application that allows you to configure and manage your 8235s and devices. Using the 8235 Management Facility you can configure, manage, and monitor the 8235s on your network, create user lists, and manage the security of your 8235s. The 8235 Management Facility is provided with all 8235s.

Figure 29 shows the 8235 Management Facility.

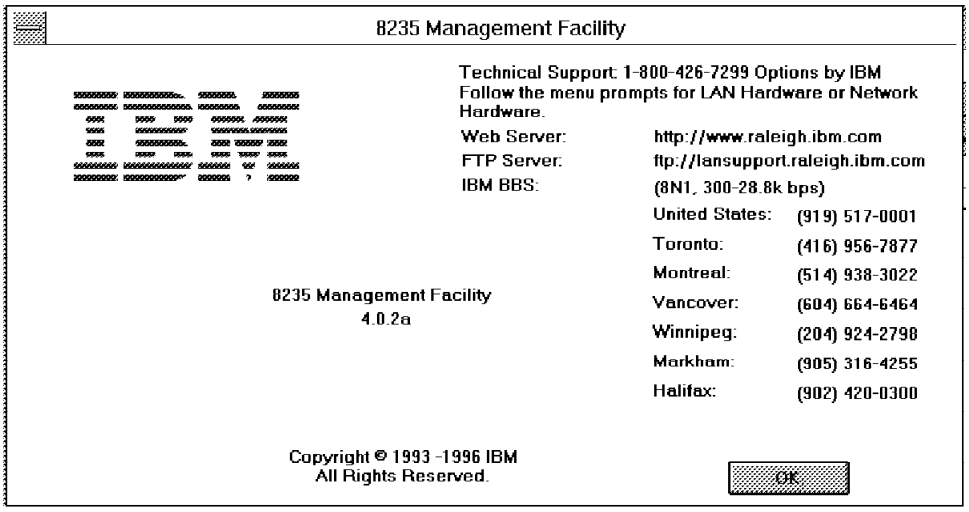


Figure 29. 8235 Management Facility

2.1.7.1 Hardware and Software Requirements

The 8235 Management Facility for Windows requires a 386, 486, or Pentium-based IBM PC or compatible workstation running Windows Version 3.1 software or Windows for Workgroups 3.11 software or higher in 386 Enhanced Mode. It is recommended that you use a 486 or Pentium PC. A mouse is required. You can also run 8235 Management Facility on a workstation running IBM WIN-OS/2 Version 3.1.

To run the 8235 Management Facility in an IPX environment, you need the Internet Packet Exchange/Sequenced Packet Exchange (IPX/SPX) ODI protocol stack from Novell, Inc. (IPXODI). The 8235 Management Facility requires the following NetWare drivers. You do not need a NetWare server on your network.

- LSL.COM Version 2.05 software or higher
- IPXODI.COM Version 2.11 software or higher
- NETX.EXE Version 3.32 software or higher or VLM Version 1.10 software or higher

Note: The 8235 Management Facility does not support the NetWare IPX.COM driver.

To run the 8235 Management Facility in an IP environment, you need a supported Winsock-compatible Internet Protocol (IP) stack. TCP/IP stacks from IBM (IBM TCP/IP for DOS Version 2.1.1), Novell, Inc. (NetWare Client Version 1.1 and LAN WorkPlace Version 4.2), and FTP (Version 3.0) are supported for use with the 8235 Management Facility over IP.

2.1.7.2 Supported Remote Access Servers

The 8235 Management Facility supports management of the following 8235s running the specified 8235 Management Facility software versions:

- 8235 Models 011, 021, 031, and 051 Versions 2.x-4.0
- 8235 Models 012, 022, 032, and 052 Versions 2.x-4.0
- 8235/T
- 8235/E

2.1.7.3 Using the 8235 Management Facility over IP

To use the Management Facility to manage the 8235s that are installed on a network, you need to install the 8235 Management Facility on a workstation that is running Windows and is using IPX or IP protocol.

The 8235 Management Facility Installation (IPX) and the 8235 Management Facility on an IP network are described, step-by-step, in *IBM 8235 Dial-in Access to LANs Server - Concepts and Experiences*, SG24-4816-00.

The Management Facility runs over one protocol stack at a time. In the Management Facility, select either the TCP/IP or IPX protocol (IPX being the default).

As we are talking about Internet service providers, we will show a basic configuration using IP protocol for TCP/IP applications, such as Internet applications.

The user interface for the Management Facility over TCP/IP is basically identical to that of IPX, except for device discovery. To start the 8235 Management Facility, you need to click twice on the **IBM 8235 Management Facility** icon in the IBM 8235 Program Group, as shown in Figure 30.

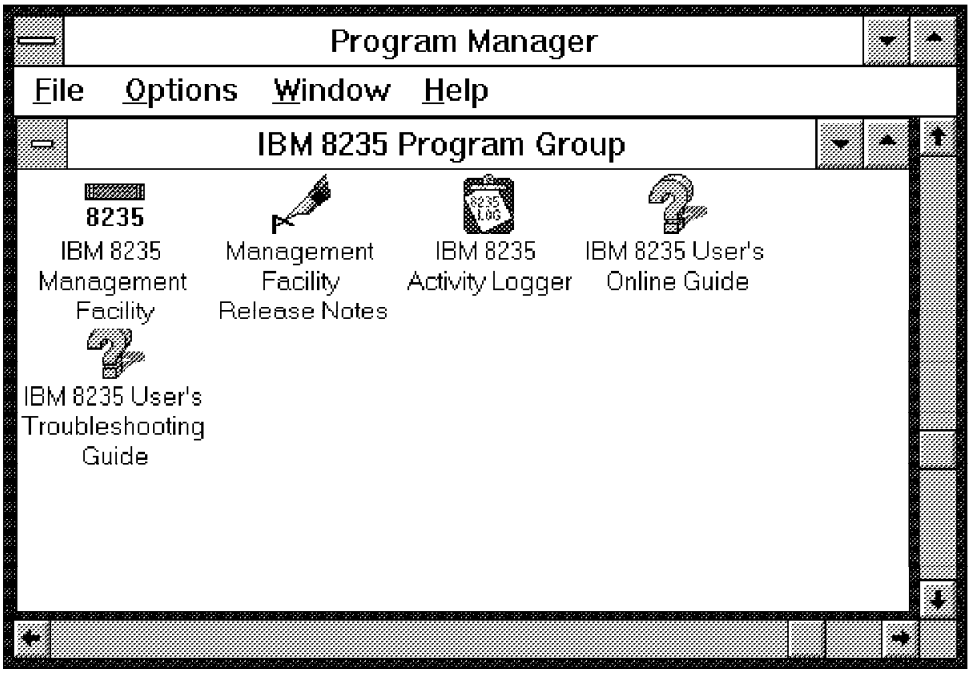


Figure 30. IBM 8235 Program Group

When you first start the Management Facility over IP, the device list is empty. IP device discovery occurs only when you select **Discover Devices** from the Devices menu. Device discovery will find IP devices on the local Ethernet or token-ring segment only. See Figure 31 on page 48.

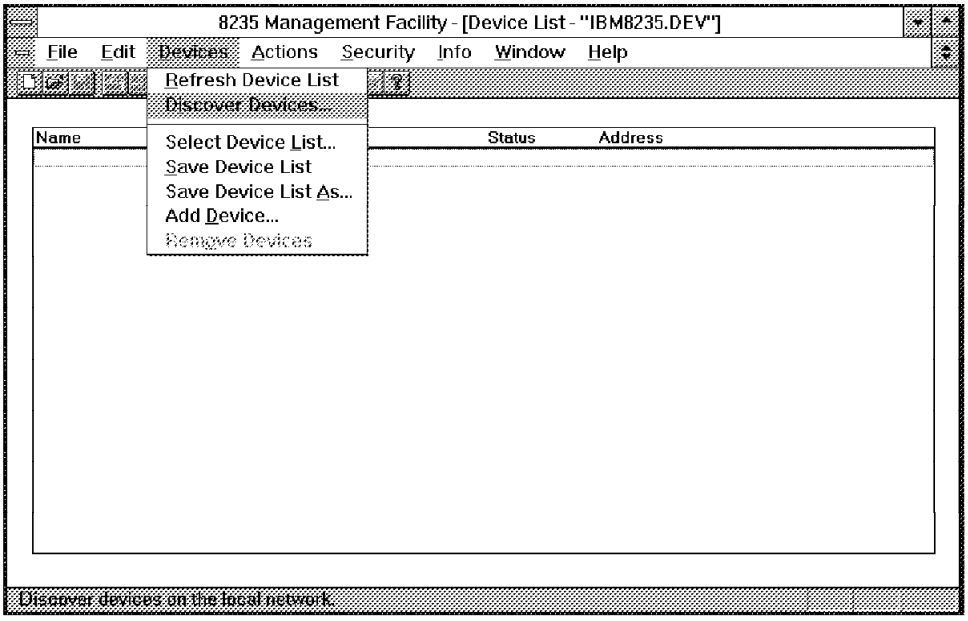


Figure 31. IBM 8235 Discover Devices on the Local Network

1. Downloading of VROM and Image Files to the 8235

 Management Facility over TCP/IP supports two types of software download to the 8235:

 - Clear and Download

 Sends VROM and Image files to a selected 8235. Refer to ″Downloading an Image and VROM Files to an 8235″ in the *8235 Management Facility User's Online Guide*.

 - Auto-Download

 With IP auto-downloading, the 8235 Management Facility automatically sends an IP address to any newly installed 8235 on your LAN. The 8235 then uses TFTP to automatically retrieve VROM and Image files. You can also download VROM and Image files to 8235s that have been pin-reset.

 You can completely manage your 8235s in an IP environment. Using Clear and Download, you can update software versions on your 8235s as well as use the commands with 8235s that already have an IP address assigned.

 Auto-Download allows you to assign an IP address to an 8235, and then download VROM and Image files. This feature allows you to quickly get new 8235s up and running as well as to upgrade existing 8235s currently installed on your network. To begin auto-download, select **Begin IP Auto-Download** from the Edit menu as shown in Figure 32 on page 49.

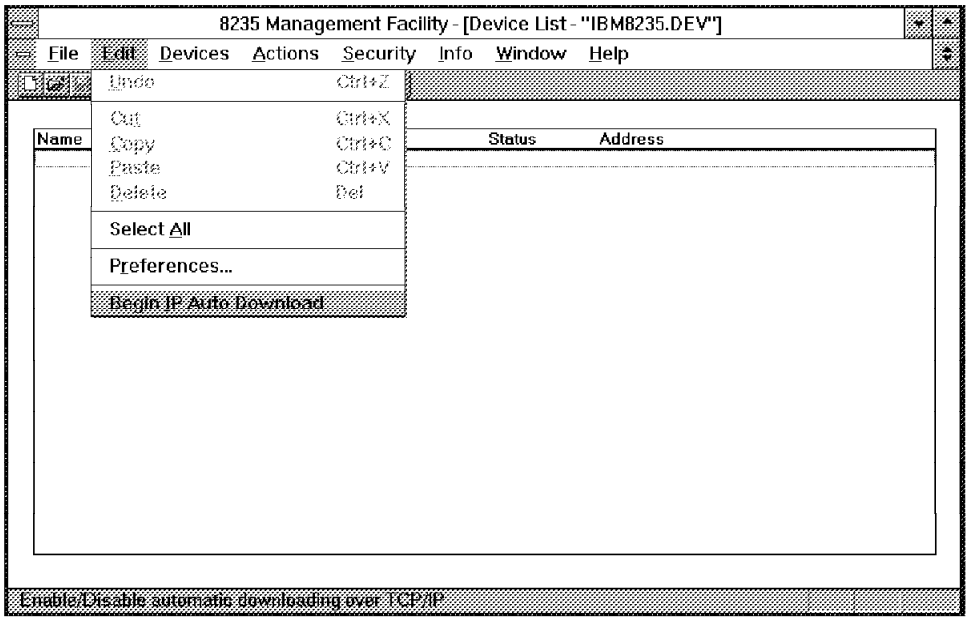


Figure 32. IBM 8235 Enable/Disable Automatic Downloading over TCP/IP

Refer to *IBM 8235 Dial-in Access to LANs Server*, SG24-4816-00 for additional information on enabling IP automatic downloading, discovering IP devices, and tips on TCP/IP.

2. Adding Devices to an IP Device List File

 The Add Devices option allows you to enter an 8235 address or IP host name. This menu option should be used to add an 8235 to the active IP device list file (see Figure 33 and Figure 34 on page 50).

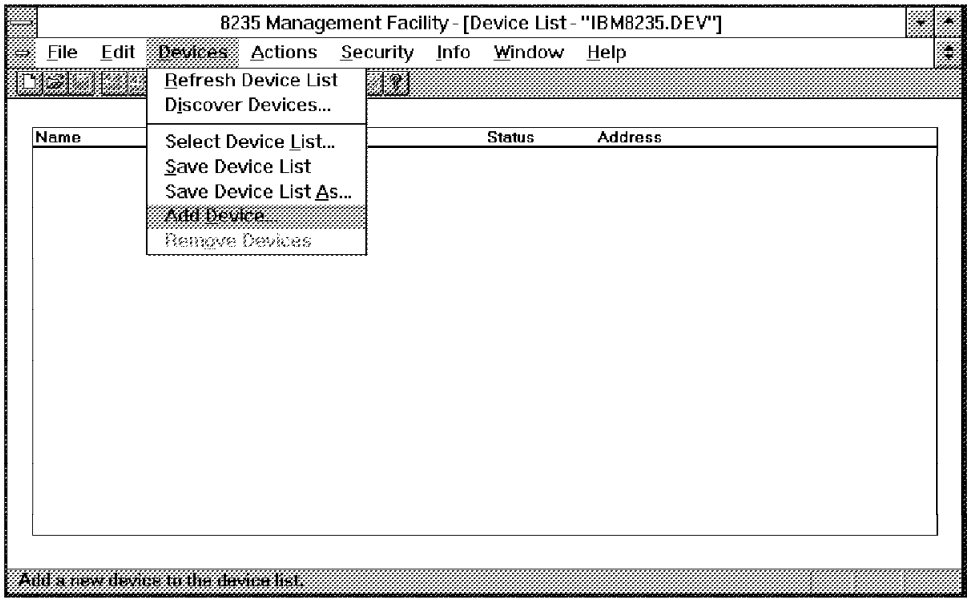


Figure 33. Add a New Device to the Device List

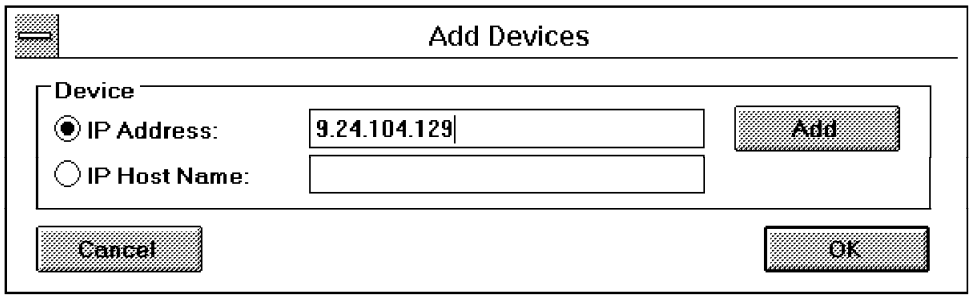


Figure 34. Add Devices Window

3. IP Device List Window

 The Device List window appears when you launch the 8235 Management Facility.

 Use the IP Devices List Window to select IP devices to configure or manage. The IP Device List window appears when you select IP in the Management Protocols page on the Preferences window, as shown in Figure 35.

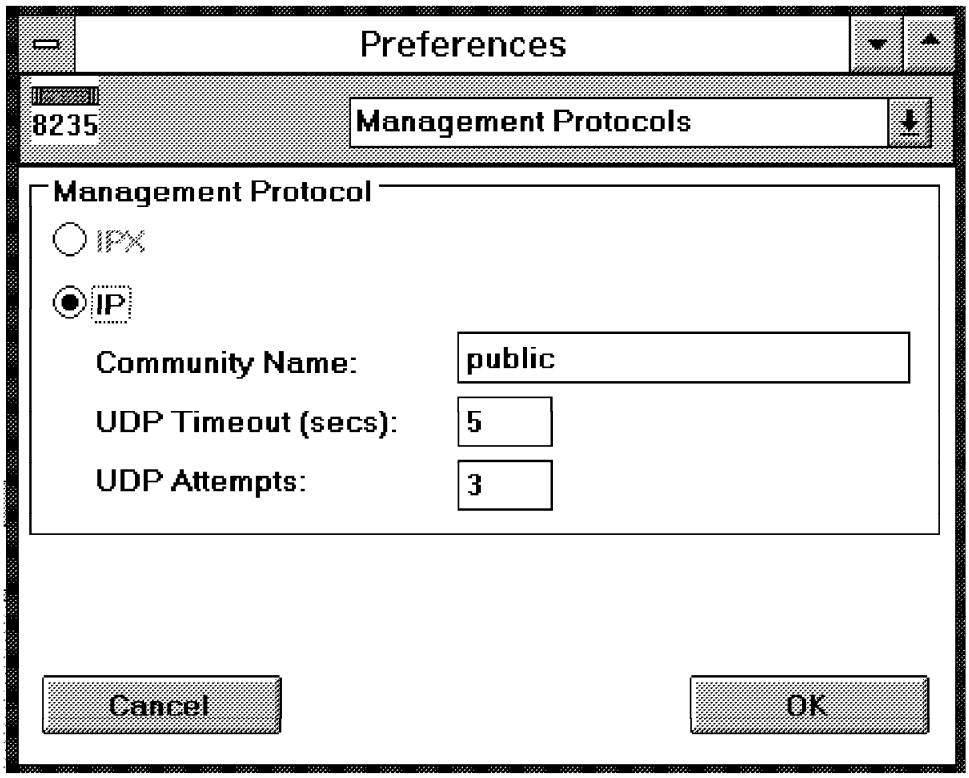


Figure 35. Management Protocols on the Preferences Window

The IP Device List window includes one list: the Device List. As we have already seen, the first time you display the IP Device window, the Device List is blank. To populate the Device List, use the Discover Devices command in the Device menu and copy one or more 8235s that you want to add to the Device List.

You can also add a device to the Device List by choosing **Add Device** from the Device menu and entering the IP address or host name of the 8235 you want to add.

Figure 36 on page 51 shows the IP Device List.

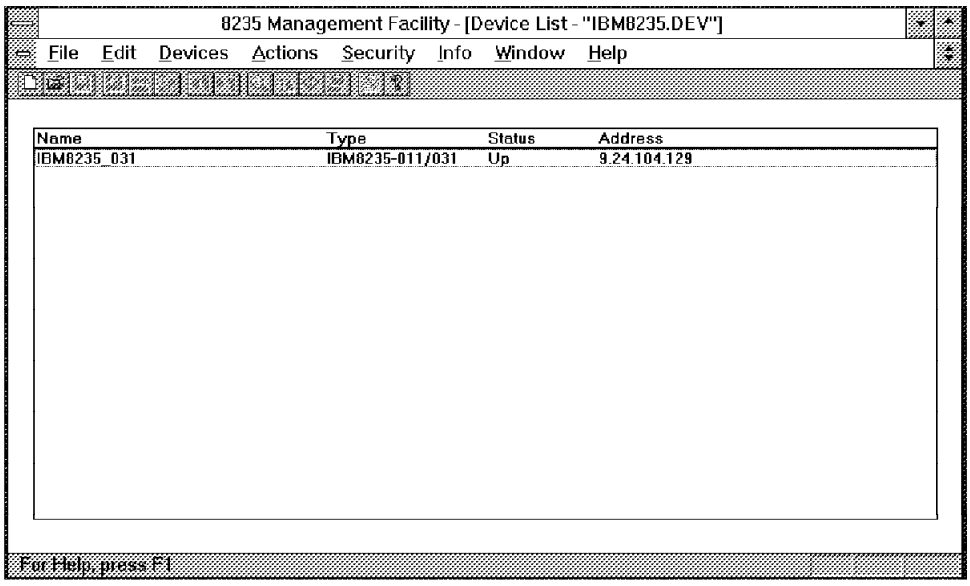


Figure 36. 8235 Management Facility - Device List

Select one or more 8235s that you wish to configure or manage and click twice on it. You will be asked to enter the Administrator Password (see Figure 37).

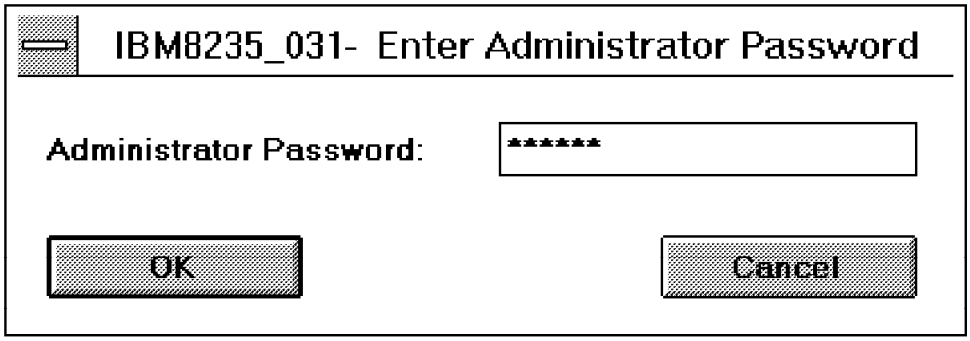


Figure 37. Enter Administrator Password Window

You should assign administrator passwords to 8235s to protect them against unauthorized access. After your identification you will have access to the Configuration window.

Use the Configuration window to edit the 8235 parameters. The Configuration window includes many pages of configuration information. To move to the next page, click on the **Configure** drop-down list at the top of the Configuration window. You have the following configuration pages:

- General Configuration page
- Ports Configuration page
- Ports: Phone Numbers Configuration page
- Virtual Connections Configuration page
- IP General Configuration page
- IP Addresses Configuration page

- IP Static Routes Configuration page
- IPX (NetWare) Configuration page
- LAN-to-LAN Sites Configuration page
- Security Configuration page
- SNMP Configuration page
- Logging Configuration page
- Bridging Configuration page
- Additional Configuration page

We show the configuration pages we need to configure the 8235 to use the IP protocol for TCP/IP applications.

4. General Configuration Page

Use the General Configuration page to edit the device name, protocols, functions, time-outs, compression, and PPP Multilink Protocol parameters. Choose **General** from the Configure drop-down list on the Configuration window to access the General Configuration page (see Figure 38).

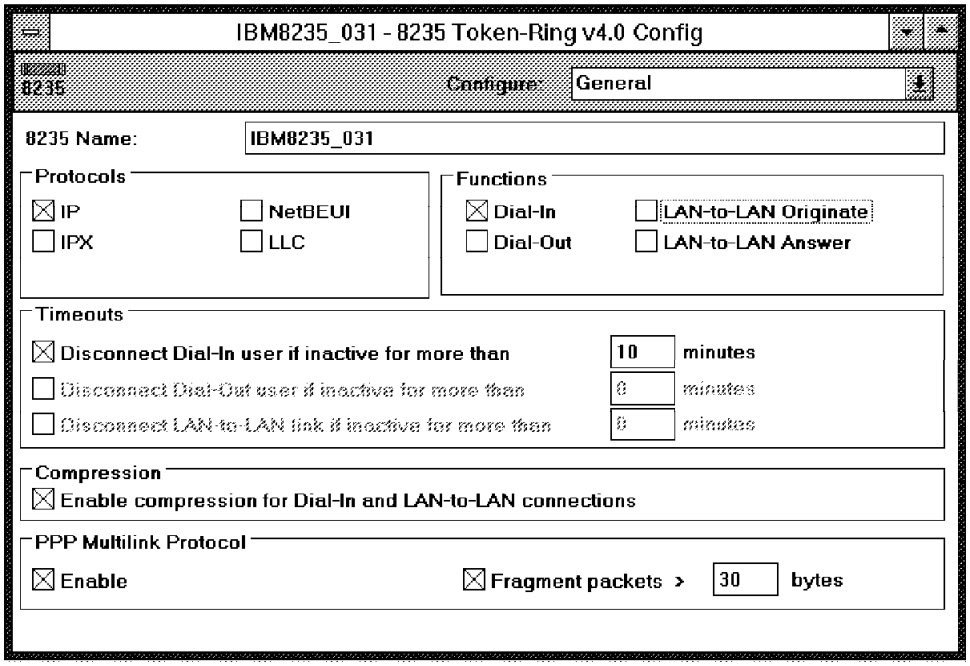


Figure 38. General Configuration Page

- Protocol Area

 Determines the protocol allowed to an 8235. The default is to enable all protocols. You must enable the protocol on this page before you can configure specific parameters on the IP, IPX, and AppleTalk Configuration pages. Select **IP** protocol.

- Functions

 Determines which functions an 8235 supports. The default is to disable all functions. Select **Dial-In** to allow users to dial in to the 8235 using one of the selected protocols.

- Timeouts

The Disconnect Dial-In User check box enables the 8235 Management Facility to disconnect inactive dial-in users after the number of minutes specified in the Minutes field. This box is selected by default. Keep it selected and enter a value from 1 to 999 minutes or deselect it.

- Compression

 Enables compression for PPP Dial-In and LAN-to-LAN connections. This check box is selected by default. If the Dial-In or LAN-to-LAN client also has data compression enabled, selecting this check box can improve the speed of dial-in connections. (If either an 8235 or the client does not have data compression enabled, this setting is ignored.)

- PPP Multilink Protocol

 Enables PPP Multilink Protocol in this device, allowing channel aggregation for dial-In and LAN-to-LAN connections.

 The Fragment Packets check box enables fragmentation of the data being transmitted via the PPP Multilink Protocol. This allows the data to be fragmented when the data packet size exceeds the number of bytes specified in the Bytes field. Fragmentation enhances load balancing across the connection links and reduces transit delay.

5. Ports Configuration Page

 Use the Ports Configuration page to select a port or channel to configure and to view a summary of port configuration settings. Choose **Ports** from the Configure drop-down list on the Configuration window to access this page (see Figure 39).

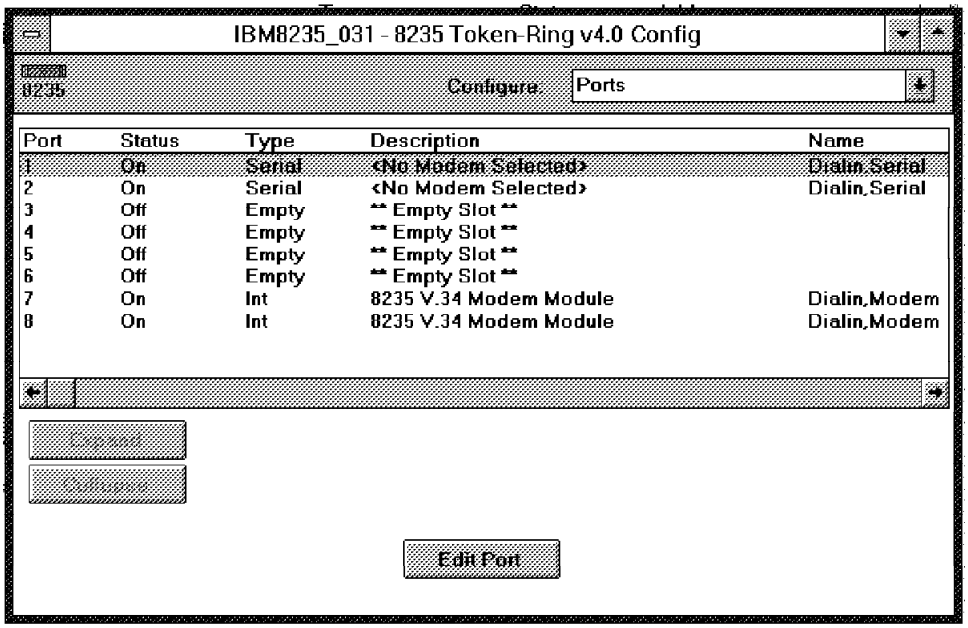


Figure 39. Ports Configuration Page

Double-click on the port or channel to configure and view the Internal Modem Module Port Configuration dialog box (see Figure 40 on page 54).

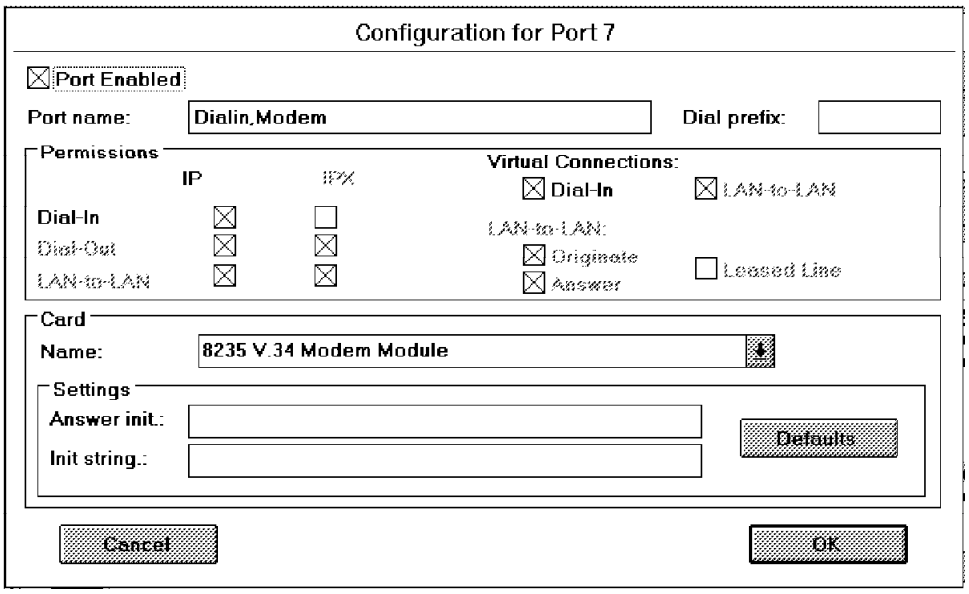


Figure 40. Internal Modem Module Port Configuration Dialog Box

- Port Enabled Check Box

 Enables the port for use. The default for this option is enabled. Even after the port is enabled, you cannot use it unless you also enable functions and protocols for the port in the Permissions area, and the corresponding protocols and functions are enabled on the General Configuration page.

- Port Name Field

 Identifies the port in the LAN-to-LAN sites and the Dial-Out Chooser windows. For detailed information, refer to *IBM 8235 Dial-In Access to LANs Server - Concepts and Experiences*, SG24-4816-00.

- Dial Prefix Field

 Dial prefix information is only used when an 8235 originates a call (either for dialback, dial-out, or originating LAN-to-LAN connections).

- Permissions Area

 Enables dial-in, dial-out, and LAN-to-LAN functions for AppleTalk, IP, and IPX protocols for the selected port or channel. The check boxes in this area are enabled only if the appropriate function or protocol has been activated in the General Configuration page. This area also enables virtual connections for Dial-In and LAN-to-LAN connections via this port or channel. The default for all permissions check boxes is enabled. Select the **Dial-in** function for the IP protocol.

 A virtual connection is a standard connection that has been enhanced to temporarily bring down the link when no meaningful data is transmitted for a specified period of time. Meaningful data includes specific requests to access or transmit information via the connection. Data that is not considered meaningful includes routine network maintenance packets.

 A virtual connection supports IP and IPX LAN-to-LAN and workstation single user, dial-in virtual connections for reduced connect-time costs and increased ease of use and management. With virtual connections the

physical connection is brought up on demand; the connection is there when you need it, and not when you do not need it.

A virtual connection can resume on a port other than the port on which the original connection was made. This is called a *floating virtual connection*. This feature eliminates the need to dedicate a particular port to each virtual connection. It also allows you to configure an 8235 for more virtual connections than the number of available ports or channels. It is possible to configure up to 200 virtual connections.

Virtual connections are ideal for ISDN connections that have quick connection times. With ISDN, resumption of dial-in and LAN-to-LAN virtual connections will be transparent to the end user. With analog dial-up connections, it could take up to 30 seconds to resume suspended virtual connections. This concept has meaning if a high-speed communication line such as T1, E1, or ISDN is being attached to the 8235. It must be an 8235 Model I40 DIALs Switch.

Note

The IBM 8235 Model I40 DIALs Switch is an enterprise-level device that attaches to one LAN (the current release supports Ethernet only) and several high-speed communication lines such as E1, T1, and primary rate ISDN (PRI) interfaces. Unlike the other 8235 models it does not directly attach to analog lines (except for its out-band management ports) or basic rate ISDN lines. However, it accepts calls from clients being attached to those lines that are being directed to its high-speed line interface by the public carrier.

For additional information, refer to *IBM 8235 Dial-In Access to LAN Servers - Concepts and Experiences*, SG24-4816-00 and to *IBM 8235 User's Online Guide - 8235 Management Facility 4.0 Release Notes*.

- Card Name Drop-Down List

 Display the list of internal devices (including modem modules) for the correct manufacturer that are stored in the MODEMS.INI file. 8235 Management Facility automatically displays the name of an internal modem module installed in the selected port.

 The 8235 Management Facility sets Answer Init and Init String fields to the values found in the MODEMS.INI file for the selected 8235.

 – Settings

 - Answer Init. Field

 Displays the command string used by an 8235 to initialize the modem when the 8235 answers a call (dial-in or LAN-to-LAN answer).

 - *Init.String Field*

 Displays the command string used by an 8235 when initiating a call (LAN-to-LAN originate or dial-out).

 Select the correct internal device (modem or modem module). The 8235 Management Facility sets Answer Init and Init String fields to the values found in the MODEMS.INI file for the selected 8235.

Figure 41 on page 56 shows the Async Serial Port Configuration dialog box. Use this port configuration dialog box to edit port configuration parameters for an 8235 using an external modem or other communications devices.

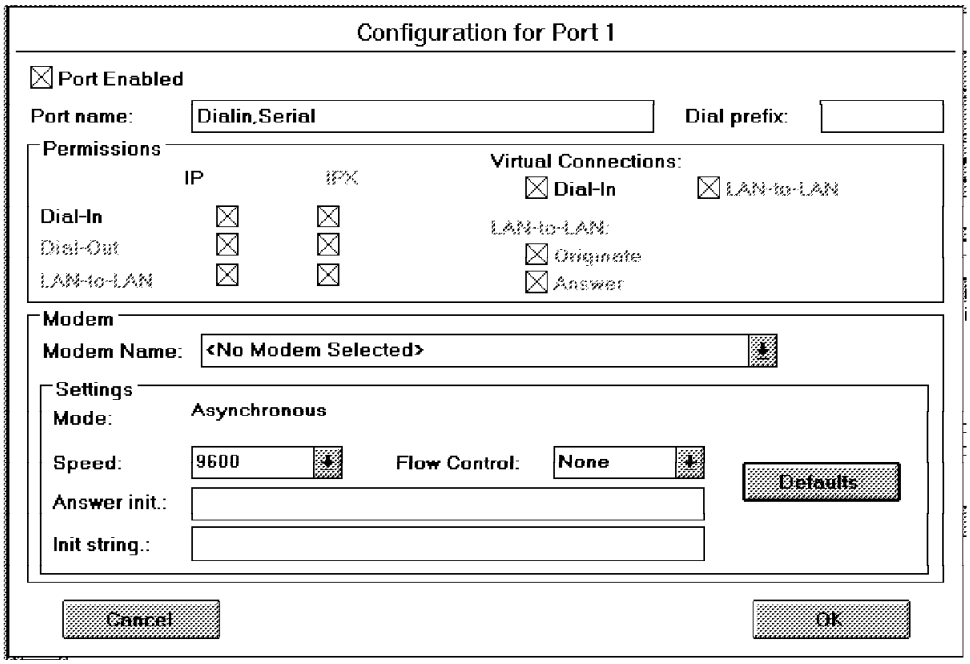


Figure 41. Async Serial Port Configuration Dialog Box

Note: This dialog box also applies to a port containing an Async Serial Module.

This dialog box is very similar to the Internal Modem Module Port Configuration dialog box.

- Permissions Area

 Select Dial-in function for IP protocol.

- Modem Name Drop-Down List

 Displays the list of modems, modem modules, terminal adapters, and ISDN adapter models and manufacturers stored in the MODEMS.INI file. For a port with an internal modem module, the 8235 Management Facility automatically selects the appropriate device from the drop-down list. For a port attached to an external device, select the name of the device (usually a modem) attached to this port. When a device is selected in the Modem Name drop-down list, the 8235 Management Facility sets the Speed, Flow Control, Answer Init., and Init. String fields to the values found in the MODEMS.INI file for the selected device.

6. Ports:Phone Numbers Configuration Page

 Figure 42 on page 57 shows the **Ports:Phone Numbers Configuration** page. This page is used to configure port and channel phone numbers. These phone numbers are used during multilink connections. Choose **Ports:Phone Numbers** from the Configure drop-down list on the Configuration window to access the Ports:Phone Numbers Configuration page. Setting up this page is not required for an 8235 used for ISP purposes.

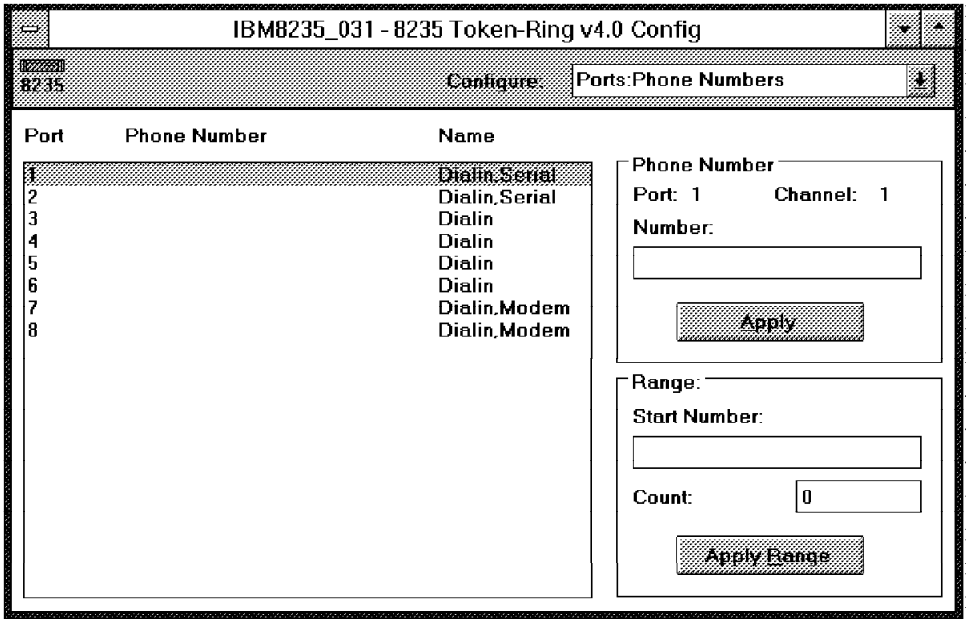


Figure 42. Ports:Phone Numbers Configuration Page

7. Virtual Connections Configuration Page

 The Virtual Connections Configuration page is used to configure dial-In and LAN-to-LAN virtual connection parameters. Choose **Virtual Connections** from the Configure drop-down list on the Configuration window to access the Virtual Connections Configuration pag (see Figure 43 on page 58).

 As we have already seen, with analog dial-up connections, it could take up to 30 seconds to resume suspended virtual connections. A high-speed communication line such as T1, E1, or ISDN, attached to an 8235, is recommended when Virtual Connections are required.

 For analog dial-up connections, do not select the Enable Virtual Connection check box on the Virtual Connections Configuration page. For T1, E1, or ISDN lines, select the **Enable Virtual Connection** check box. Before configuring virtual connections using this page, the IP protocol and Dial-In functions must be enabled on the General Configuration page.

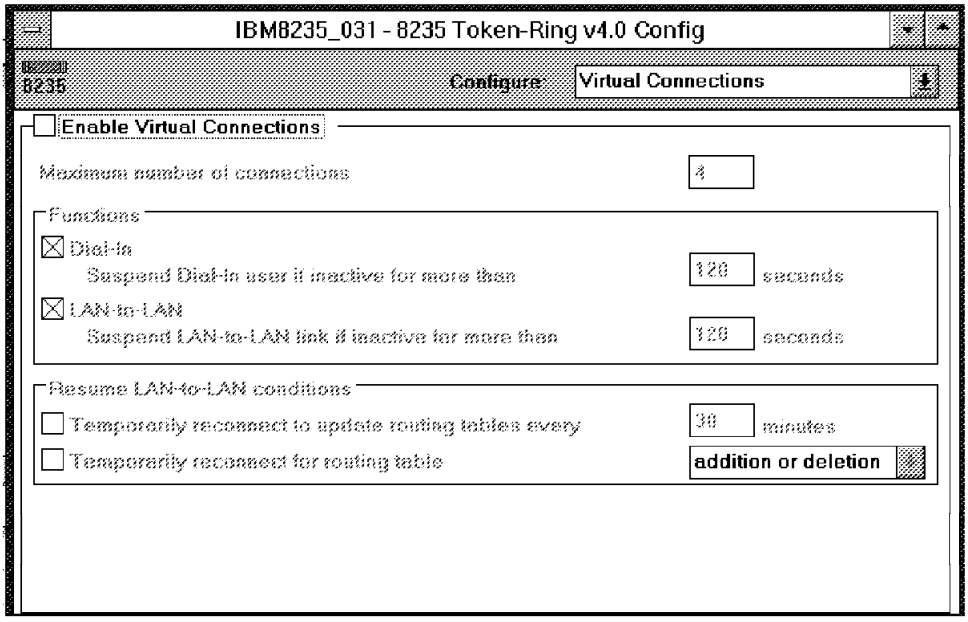


Figure 43. Virtual Connections Configuration Page

8. IP General Configuration Page

 Figure 44 shows the IP General Configuration page.

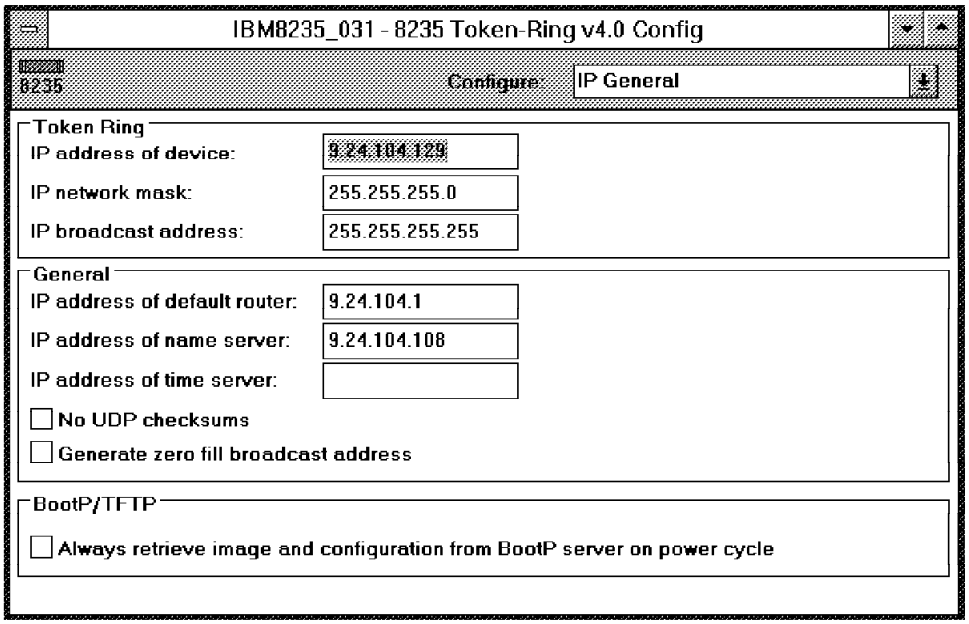


Figure 44. IP General Configuration Page

Use the IP General Configuration page to configure the Internet Protocol (IP) addresses and parameters for an 8235. Choose **IP General** from the Configure drop-down list on the Configuration window to access this page.

- IP Address of Device Field

 Sets the device's IP address, which identifies the host on the IP network. The IP address consists of a network number, which is the same for

every host on the network, and a host number, which must be unique for each host on a network.

- IP Network Mask Field

 Indicates which portions of an IP address refer to the network and which portions refer to the host. The IP network mask is also referred to as the subnet mask.

- IP Broadcast Address Field

 Sets the address used for transmitting packets that should be received and processed by all of the hosts on a given network segment.

- IP Address of Default Router Field

 Sets the IP address of a default router to which IP packets destined for remote IP hosts are forwarded by an 8235.

- IP Address of Name Server Field

 Sets the IP address of a name server host on the local IP network that translates host names into addresses using the domain name server protocol.

For additional information about the IP General Configuration page, refer to *IBM 8235 Dial-In Access to LAN Servers - Concepts and Experiences*, SG24-4816-00.

9. IP Addresses Configuration Page

 Figure 45 shows the IP Addresses Configuration page.

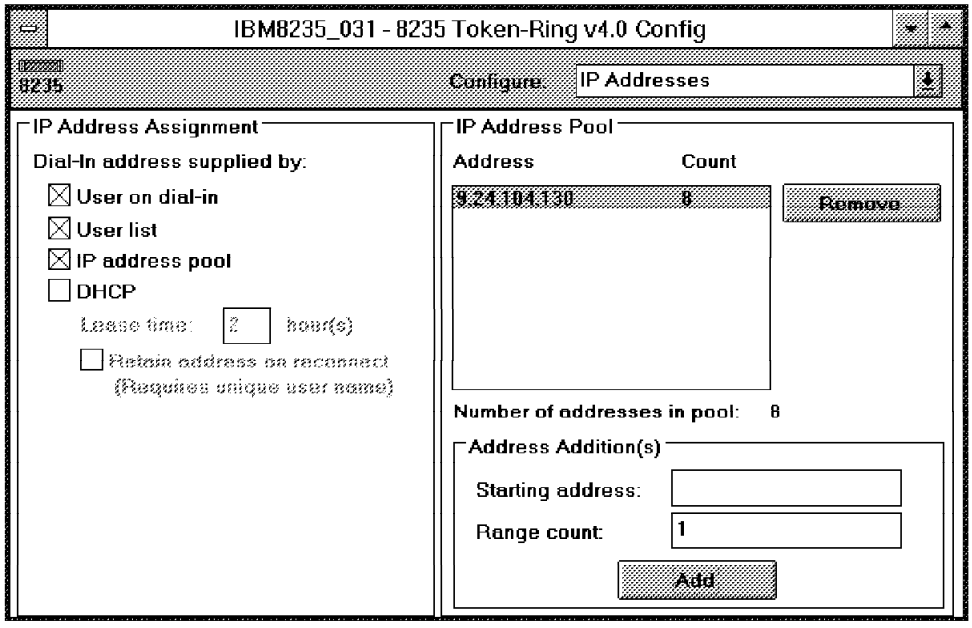


Figure 45. IP Addresses Configuration Page

Use the IP Addresses Configuration page to assign the Internet Protocol (IP) addresses for dial-in users and to configure an IP address pool. Choose **IP Addresses** from the Configure drop-down list on the Configuration window to access this page.

- IP Address Assignment Area

Allows the dial-in user, user list, port, or Dynamic Host Configuration Protocol (DHCP) to supply the IP address for a dial-in user. The address might be changed dynamically: an 8235 does not have to be restarted for a change in the IP address policy to take effect. The precedence of the address sources is: user on dial-in, user list, port, then DHCP. You can assign more than one address source.

– User on Dial-In Check Box

 Enables a user-specified IP address on dial-in. When dialing into a network, users can enter an IP address of their choice. (This address must be valid for the network.)

– User List Check Box

 Enables the user list to supply the IP address.

– IP Address Pool Check Box

 Enables the IP address pool to supply the IP address. When this check box is active, the dial-in user is assigned the first available IP address from the IP address pool upon connection.

– DHCP Check Box

 Enables a Dynamic Host Configuration Protocol (DHCP) server on the network to dynamically assign the IP address. This option does not work for LAN-to-LAN connections. Selecting this check box enables the IP Address Lease Time field and IP Address Retained on Reconnect check box.

 - Lease Time Field

 Sets the DHCP IP address lease time in hours. Use a short lease time (1-3 hours) to conserve the IP address on the network. Use a long lease time (up to 48 hours) to increase the chance of the user getting the same address when reconnecting. The default value is 2 hours.

 - Retain Address on Reconnect Check Box

 Enables dial-in users to retain their IP addresses between dial-in sessions. This option requires that dial-in users have unique user names.

• IP Address Pool Area

 Allows you to configure the IP address pool for an 8235.

 – IP Address Pool List

 Lists the IP addresses that can be assigned to dial-in users upon connection.

 – Address Addition(s) Area

 Allows you to add IP addresses to the IP Address Pool list.

 - Starting Address Field

 Displays the IP address for the selected entry in the IP Address Pool list.

 To add several consecutive IP addresses, enter the starting IP address in this field and use the Range Count field to specify the number of addresses in the range.

- Range Count Field

 Sets the number of IP addresses will that be added to the IP Address Pool list.

 The default for this field is 1. To add more than one consecutive IP address, the 8235 Management Facility increments the starting address by 1 for each address in the series. For example, if the starting IP address is 140.124.250.145, and we have a range count of 3, the 8235 Management Facility allocates the IP addresses 140.124.250.145, 140.124.250.146 and 140.124.250.147 for an 8235.

10. IP Static Routes Configuration Page

 Use the IP Static Routes Configuration page to configure a set of permanent routes in an 8235. Choose IP Static Routes from the Configure drop-down list on the Configuration window to access this page (see Figure 46).

 Static routes are useful when selecting a preferred route to a remote host, or on internetworks that use routing protocols other than RIP. Each permanently configured IP address is known as a static route.

 This page is available only when the IP protocol is enabled on the General Configuration page.

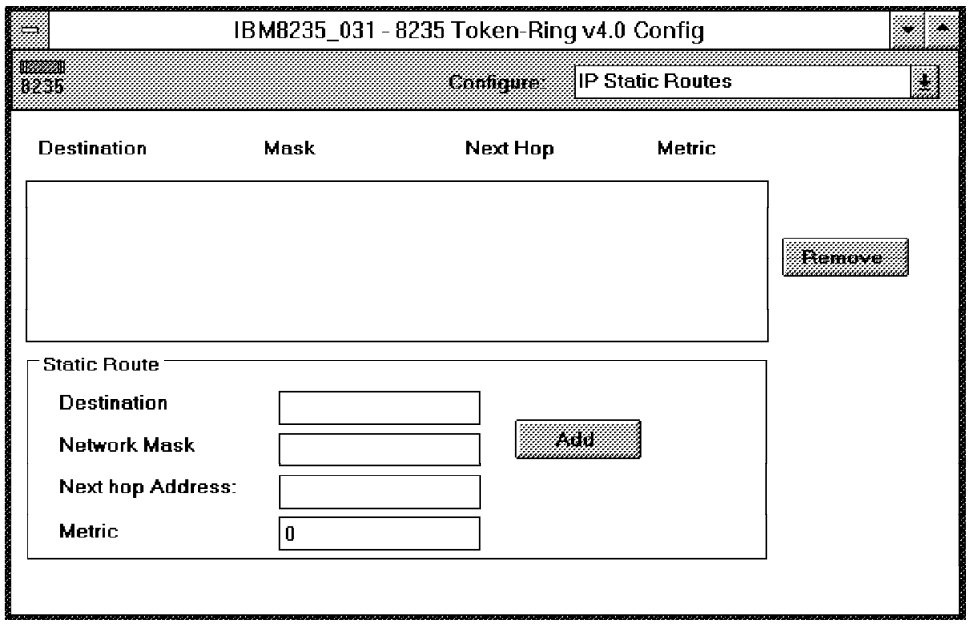


Figure 46. IP Static Routes Configuration Page

- Destination Field

 Sets the destination of the static route. The destination must be an IP address (entered in dotted decimal notation); domain names are not accepted. If the destination is a network, the node portion of the IP address is 0. If the destination is a host, the mask must be 255.255.255.255.

- Network Mask Field

 Indicates the network and subnet portion of the IP address with non-zero numbers; the node portions are shown with zeros.

- Network Hop Address Field

Sets the address of the next-hop router. The next-hop router must be on the same local network as an 8235.

- Metric Field

 Indicates the number of hops between an 8235 and the destination.

11. Security Configuration Page

 Use the Security Configuration page to configure the extended security features of an 8235. Choose **Security** from the Configure drop-down list on the Configuration window to access this page (see Figure 47).

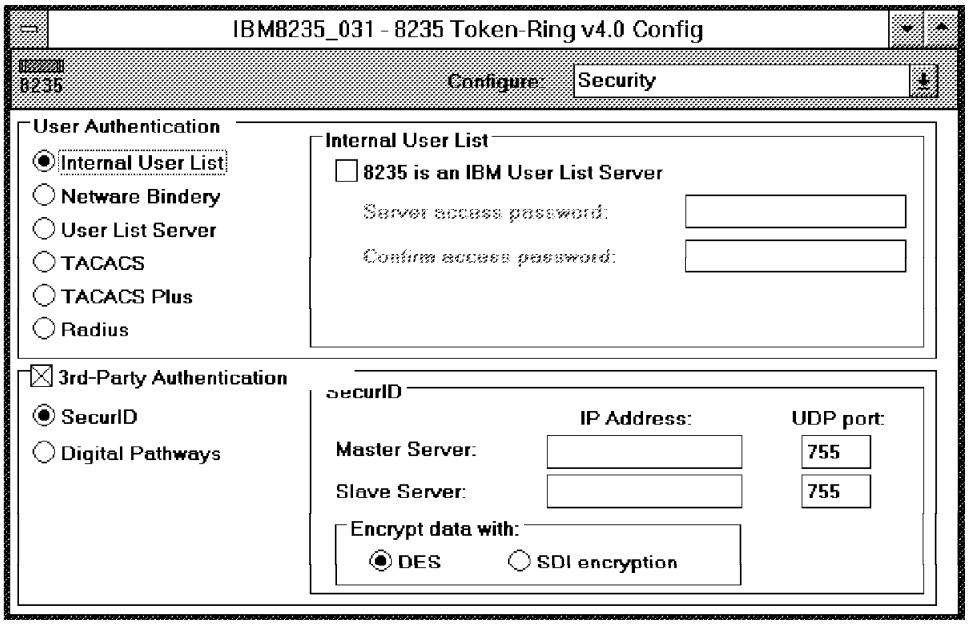


Figure 47. Security Configuration Page

Configuration of the security page is dynamic; it is not necessary to restart the device for changes to take effect. Instead, changes in the security configuration take effect on the next user authentication performed.

- User Authentication Area

 Allows you to select how primary user authentication is accomplished. The information in this area changes depending on the selected authentication method.

 – Internal User List Radio Button

 Enables an 8235 to authenticate users by verifying them against the 8235's internal user list. When this radio button is activated the Internal User List area appears (See Figure 48 on page 63).

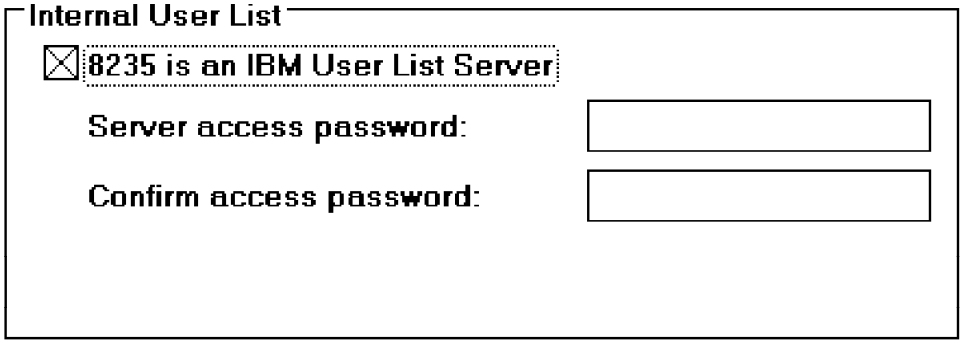


Figure 48. Internal User List Area

8235 is an IBM User List Server Check Box -
Enables the 8235 to act as a central user list server. This allows other 8235s to share this 8235's user list for user authentication.

Server Access Password Field Sets the password required to share this 8235's user list for user authentication.

Confirm Access Password Field Confirms the password.

– NetWare Bindery Radio Button

Enables the device to use the Bindery database of a NetWare Server for user authentication. When this radio button is activated the NetWare Bindery area appears (see Figure 49).

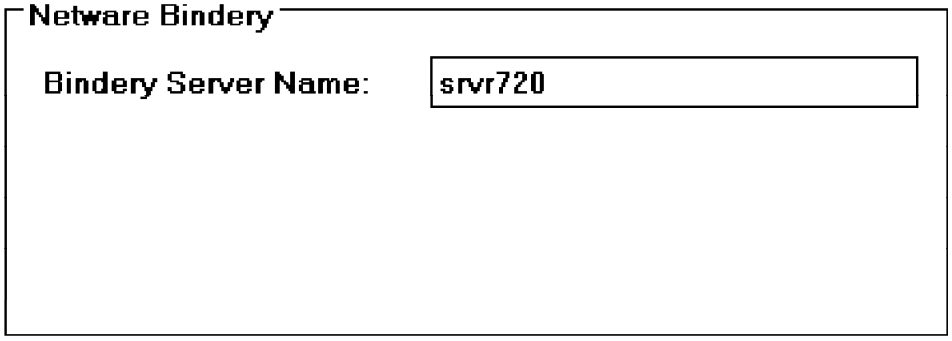


Figure 49. NetWare Bindery Area

Bindery Server Name Field Indicates the name of the main Bindery server to use.

– 8235 User List Server Radio Button

Enables an 8235 to authenticate users by reading the user list in another 8235 that is acting as an 8235 User List Server. When this radio button is activated the 8235 User List Server area appears (see Figure 50 on page 64).

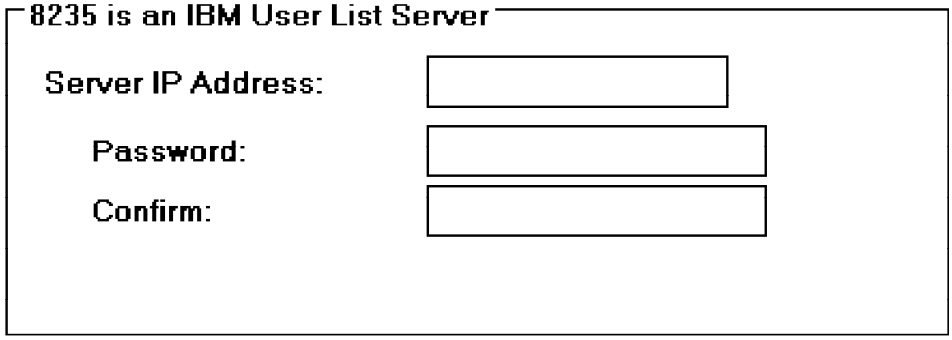


Figure 50. 8235 User List Server Area

Server IP Address Field Sets the IP address of the 8235 User List Server that the device accesses for user authentication.

Password Field Sets the password used to access the 8235 User List Server.

Confirm Field Confirms the 8235 User List Server's password.

– TACACS Radio Button

TACACS (Terminal Access Controller Access Control System) is an industry-standard security protocol. When a user attempts to gain access (such as a remote user logging in to a network), a TACACS system forwards the user name and password information to a centralized server. This server performs the necessary verification and sends a response back to the TACACS system to either allow or deny the access to the network. When this radio button is activated the TACACS area appears (see Figure 51).

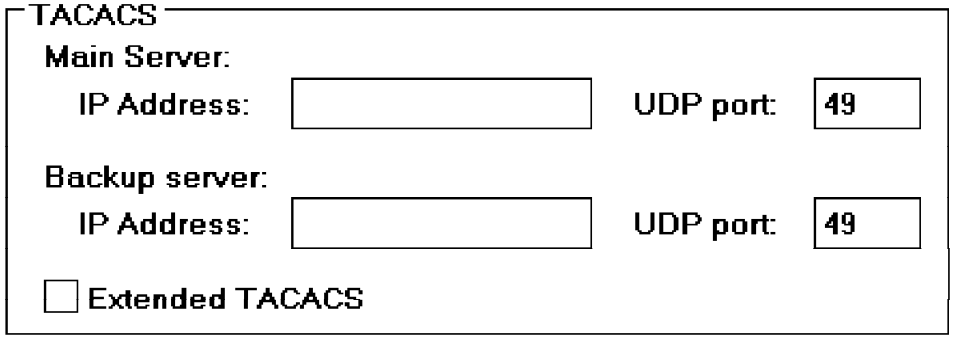


Figure 51. TACACS Area

Main Server IP Address Field Sets the IP address, in dotted-decimal notation, of the main TACACS server.

Main Server UDP Port Field Sets the new UDP port number if the original has been changed; otherwise, uses the default value of port 49.

Backup Server IP Address Field Sets the IP address, in dotted-decimal notation, of the backup TACACS server.

Backup Server UDP Port Field Sets the new UDP port number if the original has been changed; otherwise, use the default value of port 49.

– TACACS Plus Radio Button

Enables an 8235 to use Terminal Access Controller Access Control System (TACACS) Plus, an enhanced version of the TACACS security protocol, for user authentication. TACACS Plus is a security protocol used to communicate between a device and an IP authentication database. When this radio button is activated the TACACS Plus area appears. See Figure 52.

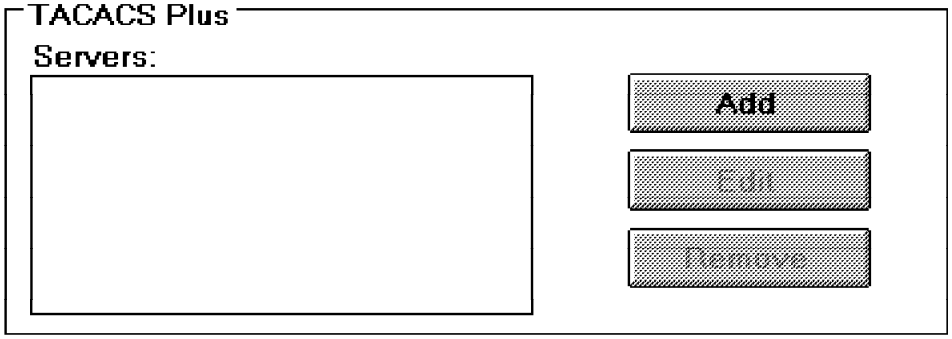


Figure 52. TACACS Plus Area

Servers List Field Lists the TACACS Plus servers on the network that an 8235 accesses for user authentication.

Add Button Displays the TACACS Plus Server dialog box, which allows you to add information for a TACACS Plus Server to the Servers list.

Use the TACACS Plus Dialog Box to add or edit information for a TACACS Plus server used for user authentication (see Figure 53).

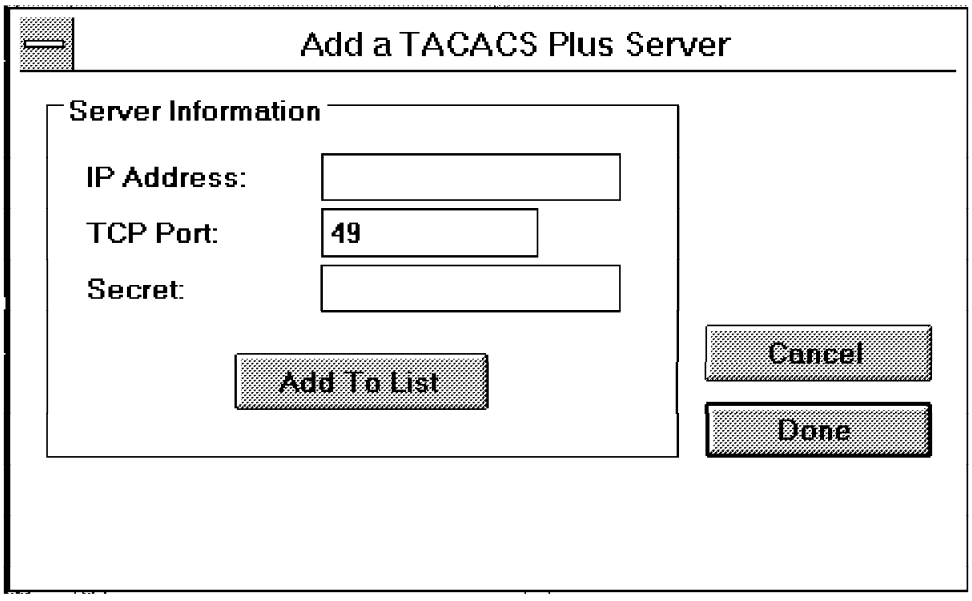


Figure 53. TACACS Plus Dialog Box

IP Address Field Sets the IP address of the TACACS Plus server.

TCP Port Field Specifies the number of the port that the TACACS Plus server uses to communicate. The default value for the TACACS Plus server TCP Port field is 49.

Secret Field Specifies the secret key used by the TACACS Plus server and an 8235 to encrypt data packets.

Add To List Button Adds the TACACS Plus server information specified in the TACACS Plus Server dialog box to the Servers list.

Done Button Saves changes and closes the TACACS Plus Server dialog box.

Edit Button Displays the TACACS Plus Server dialog box, which allows you to edit information for the selected TACACS Plus Server.

Remove Button Removes the selected TACACS Plus Server from the Servers list.

– Radius Radio Button

Enables an 8235 to access a radius server for user authentication and authorization. When this radio button is activated the Radius area appears (see Figure 54).

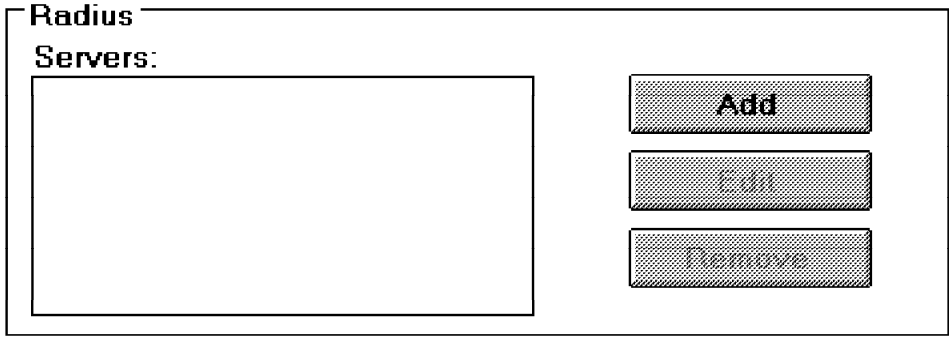


Figure 54. Radius Area

Servers List Field Lists the radius servers on the network that an 8235 accesses for user authentication. Server list entries include the server's IP address and secret. To edit the server information, double-click on the server entry. The 8235 Management Facility allows you to configure up to three radius servers.

Add Button Displays the Radius Server dialog box, which allows you to add information for a radius server to the Servers list.

Use the Radius Dialog Box to add or edit information for a radius server used for user authentication (see Figure 55 on page 67).

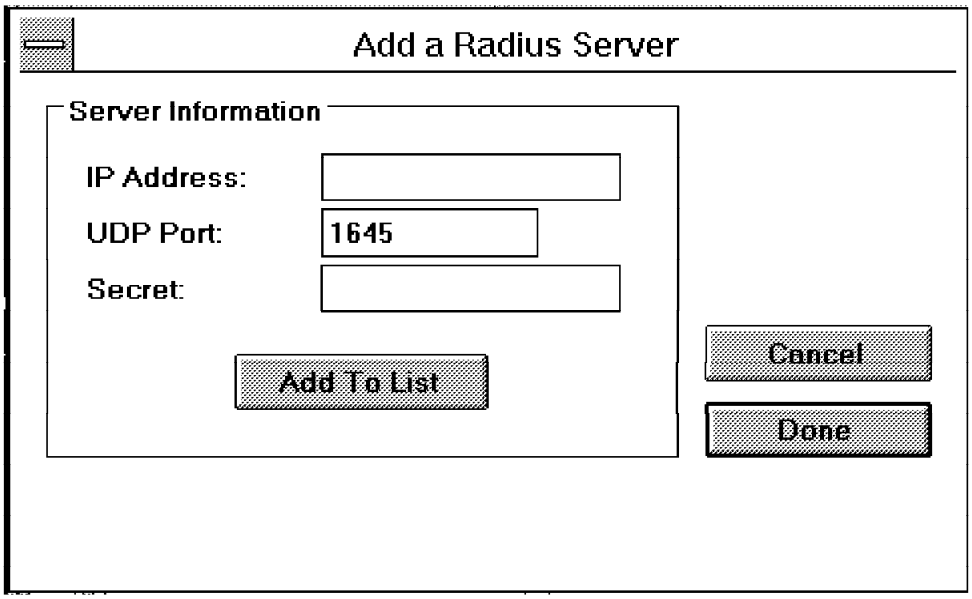


Figure 55. Radius Dialog Box

IP Address Field Sets the IP address of the radius server.

TCP Port Field Specifies the number of the port that the radius server uses to communicate. The default value for the radius server TCP Port field is 1645.

Secret Field Specifies the secret key used by the radius server and an 8235 to encrypt data packets.

Add To List Button Adds the radius server information specified in the radius server dialog box to the Servers list.

Done Button Saves changes and closes the Radius Server dialog box.

Edit Button Displays the Radius Server dialog box, which allows to edit information for the selected Radius Server.

Remove Button Removes the selected Radius Server from the Servers list.

– Third-Party Authentication Check Box

Enables third-party authentication for an 8235 in addition to the main authentication method selected in the User Authentication Area. Activating this check box enables the SecurID and Digital Pathways radio buttons.

SecurID Radio Button Enables the device to authenticate users using SecurID. When this radio button is activated the SecurID area appears (see Figure 56 on page 68).

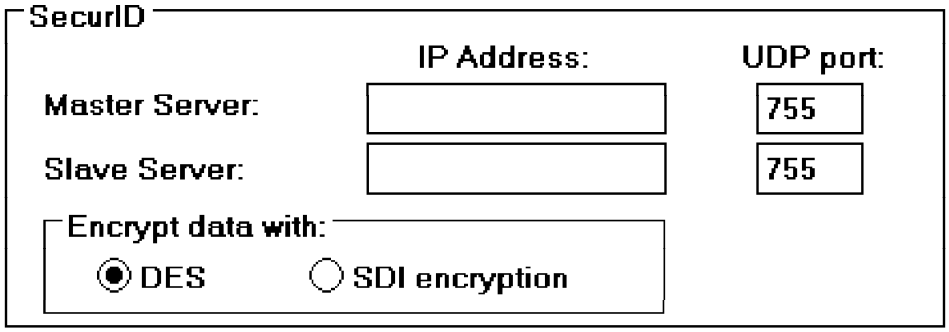


Figure 56. SecurID Area

Master Server IP Address Field Displays the IP address of the main SecurID server.

Master Server UDP Port Field Displays the UDP port number of the master SecurID server.

Slave Server IP Address Field Displays the IP address of a backup SecurID server. An 8235 accesses the slave SecurID server if the master server is unavailable.

Slave Server UDP Port Field Displays the UDP port number of the slave SecurID server.

Encrypt Data Radio Buttons Indicates the method used to encrypt data. Options include DES and Security Dynamics Inc. (SDI) encryption.

Digital Pathways Radio Button Enables the device to authenticate users using a digital pathways server. When this radio button is activated the Digital Pathways area appears (see Figure 57).

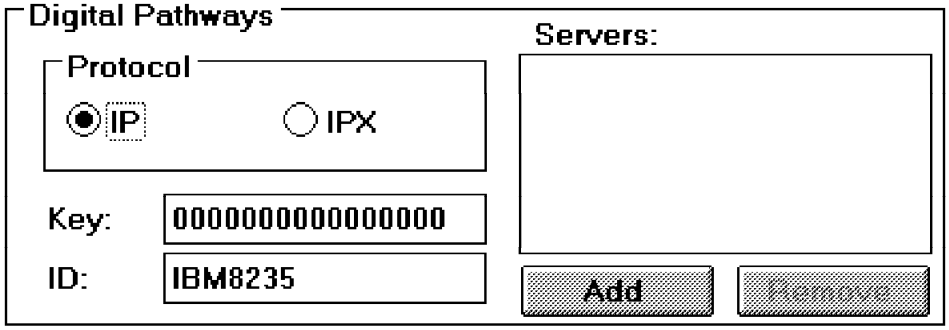


Figure 57. Digital Pathways Area

Protocol Radio Buttons Enable either IP or IPX to specify the protocol to use to connect to the Digital Pathways server. Select the **IP** radio button.

Key Field Enter the AgentKey for the 8235. This 16-digit, hexadecimal number must also be configured in the Digital Pathways server, which uses this value to authenticate the 8235 before user authentication.

ID Field Enter an alphanumeric AgentID for the 8235. This case-sensitive ID must also be configured in the Digital Pathways server, which uses this ID to authenticate the 8235 before user authentication.

Servers List Field Lists the Digital Pathways servers on the network that an 8235 accesses for user authentication. To add a server for the selected protocol, click **Add**, enter the appropriate server information and then click **Done**. For IP servers, the server's IP address and TCP port number are required (see Figure 58).

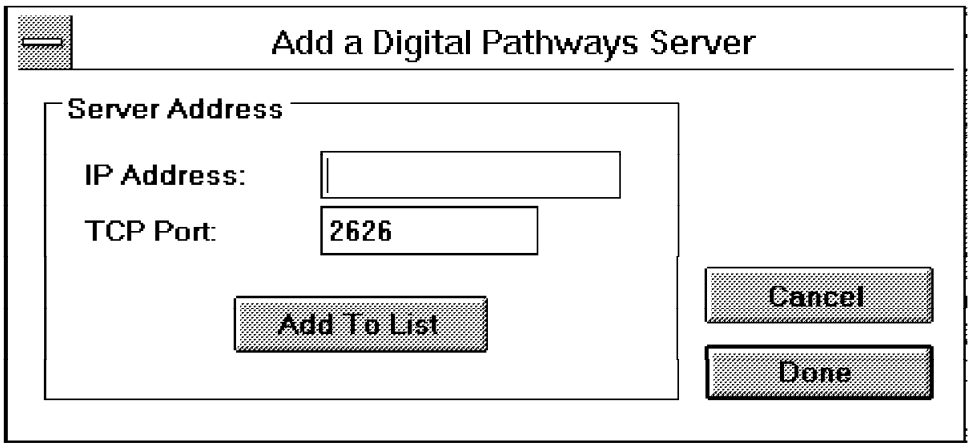


Figure 58. Digital Pathways Dialog Box

It is possible to configure a primary and a backup server for each protocol. The first server listed for a particular protocol is treated as the primary server. During user authentication, the 8235 attempts to access the first valid server listed for the selected protocol. If this attempt fails, it tries to connect to the next valid server for that protocol. If the attempt fails again, the 8235 cycles back to the first server and tries again. The 8235 continues cycling through the Server list for that protocol until it successfully connects.

For additional information about:

- Security Dynamics, refer to http://www.securid.com
- Digital Pathways, refer to http://www.digpath.com

– SNMP Configuration Page

Use the SNMP Configuration page to configure the Simple Network Management Protocol (SNMP) network management settings for an 8235. Choose **SNMP** from the Configure drop-down list on the Configuration window to access this page (see Figure 59 on page 70).

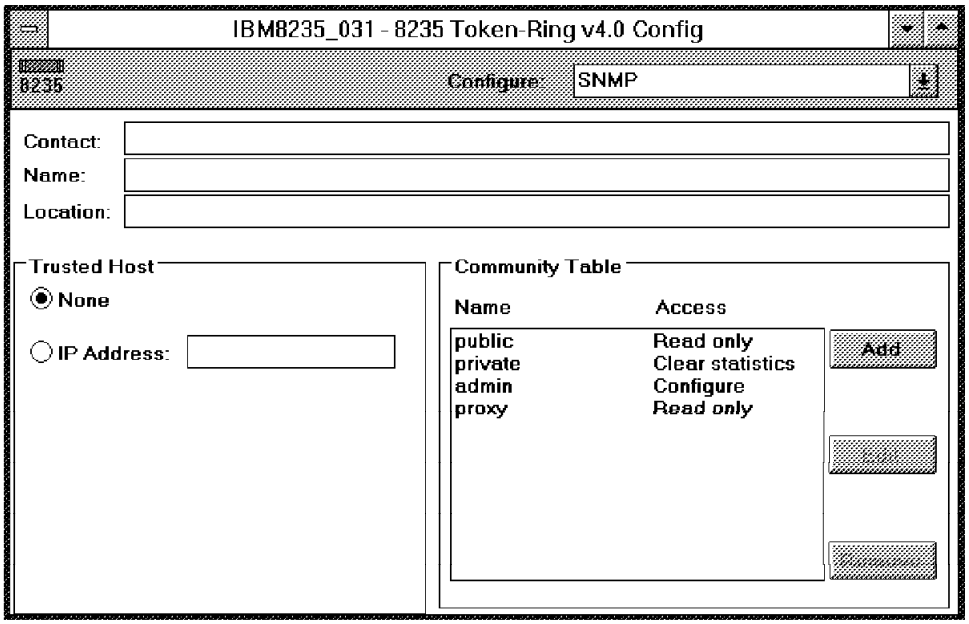


Figure 59. SNMP Configuration Page

For additional information concerning each page included in the Configuration window, refer to:

- *IBM 8235 Dial-In Access to LANs Server - Concepts and Experiences*, SG24-4816-00
- IBM 8235 User's Online Guide

- Routing Table Window

 Use the Routing Table window to view the list of networks recognized by an 8235. Select **Routing Table** from the Info menu (see Figure 60).

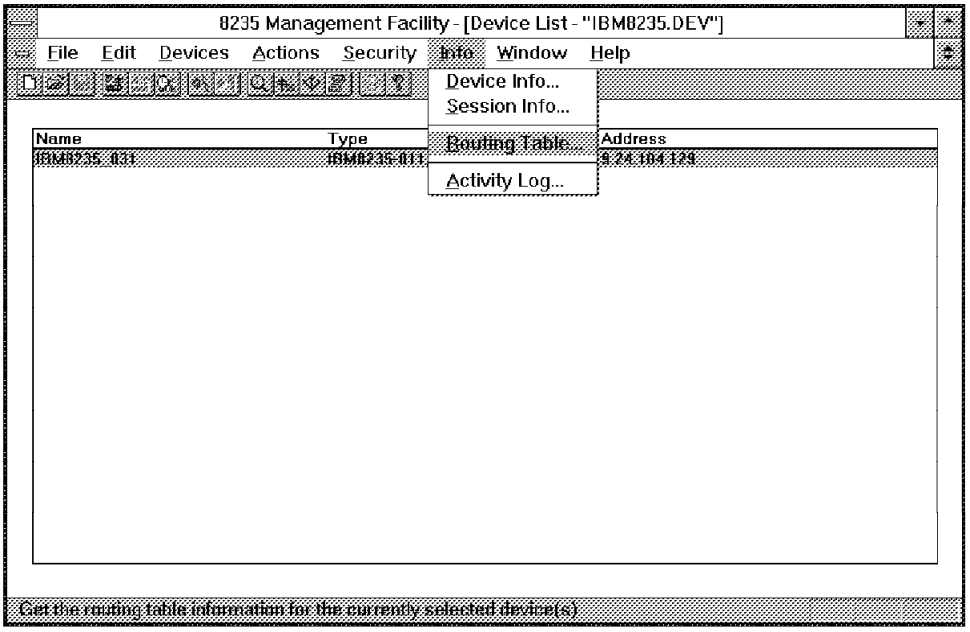


Figure 60. routing table from the Info Menu

Figure 61 on page 71 shows a Routing Table.

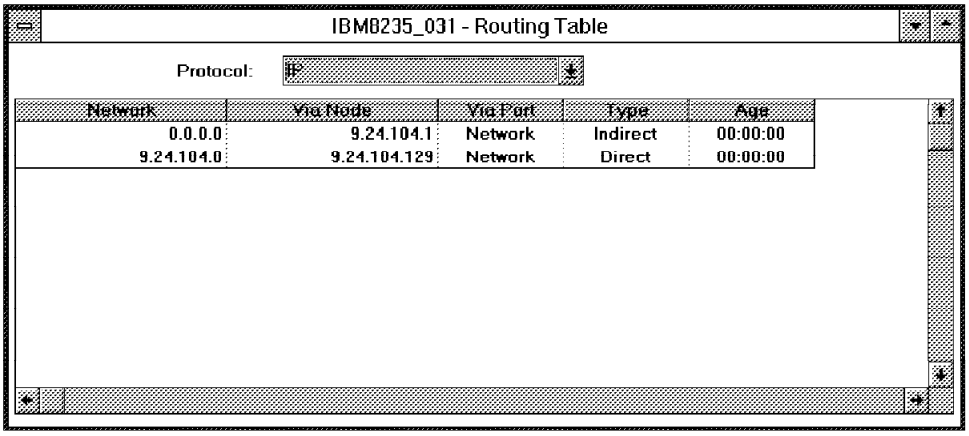


Figure 61. Routing Table

Use the IP Routes page to view IP networks recognized by an 8235. The fields are for display only.

Network Field Lists the network number.

Via Node Field Indicates the node number of the router used to forward packets to this network.

Via Port Field Indicates the 8235 port used for this route.

Type Field Indicates the IP routing protocol used.

Age Field Indicates the age of the network connection.

2.1.8 8235 Hardware

Figure 62 shows the front panel for all models of the 8235.

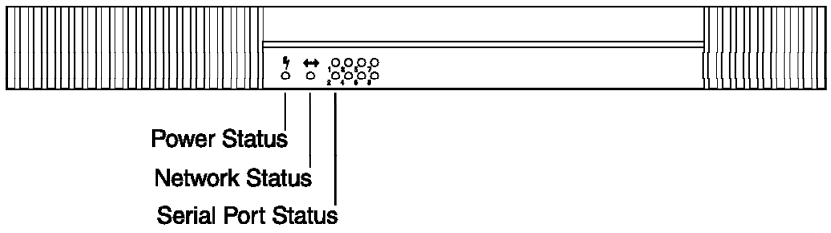


Figure 62. 8235 Front View

The front panel contains LEDs that indicate:

- Power status
- Network status
- Serial port status

Table 7 shows the meanings of the status indicator LEDs on the front panel of the 8235 in various operating modes, and Table 8 on page 72 shows the meaning of the power LED.

Table 7 (Page 1 of 2). Meanings of 8235 Network Status and Port Status LEDs

Status	Network Status LEDs	Port Status LEDs
OFF	No power or no network connection	Not in use

Table 7 (Page 2 of 2). Meanings of 8235 Network Status and Port Status LEDs		
Status	**Network Status LEDs**	**Port Status LEDs**
Green	Connected to network but idle	User connected
Green flashing (consistent)	Downloading microcode	Downloading microcode
Green flashing (inconsistent)	Connected to the network and transmitting	User connected
Green and Orange flashing	Connected to the network and transmitting with errors	
Orange flashing (consistent)	Power on self-test	Port configuration errors
Orange flashing (inconsistent)	Connected and transmitting with errors	Connected to the modem and transmitting with transmit or receive errors
Orange (solid)	8235 hardware failure	Port or 8235 hardware failure

Table 8. Meaning of 8235 Power Status LED	
Status	**Meaning**
ON	Indicates that the 8235 is powered on

2.1.8.1 LAN Connection

As mentioned earlier, the 8235 comes in two models:

- Model 1 contains a token-ring connection port.
- Model 2 has an Ethernet connection port.

The 8235 is also available as a module for the 8250 multiprotocol hub in token-ring and Ethernet models. Figure 63 shows the rear view of the token-ring Model 8235-021.

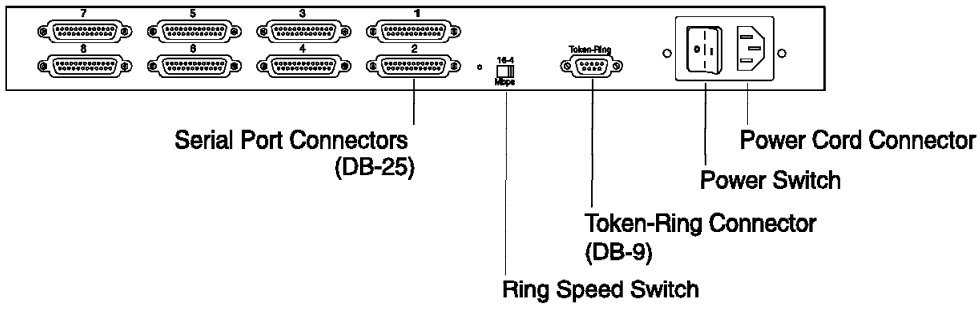


Figure 63. 8235 Model 021 Rear Panel

Figure 64 on page 73 shows the rear panel of the token-ring model 8235-031.

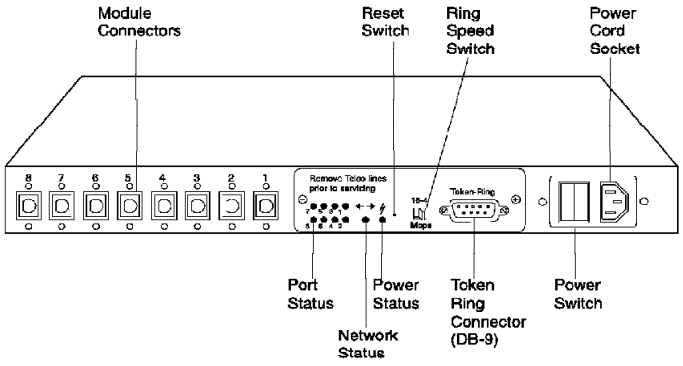


Figure 64. 8235 Model 031 Rear Panel

You make all connections on the 8235 rear panel, so the token-ring model includes one token-ring connector (DB-9) and a ring data rate switch to select the data rate of 4 or 16 Mbps.

Note

The data rate you set must match the data rate of the token-ring network. Be sure to set the power switch to Off (O) before you set the data rate.

Figure 65 shows the rear panel of the 8235 Ethernet Model 022.

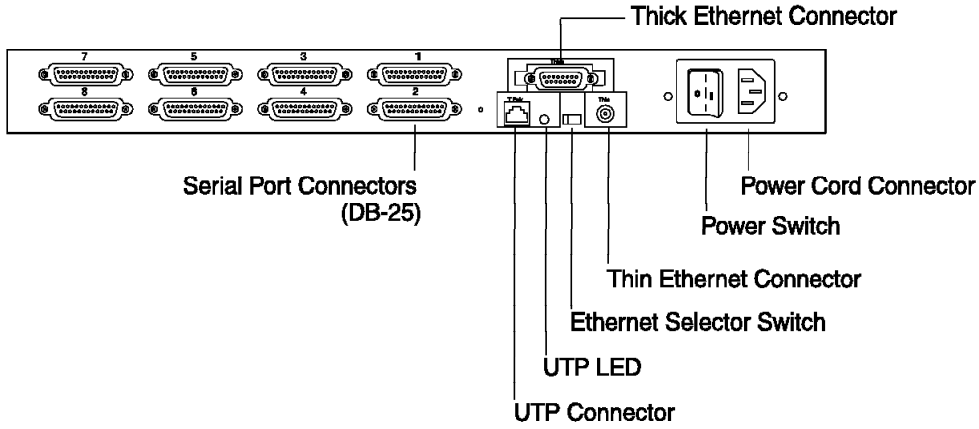


Figure 65. 8235 Model 022 Rear Panel

The 8235 Model 022 (Ethernet) provides three connectors for Ethernet: AUI (Thick Ethernet), BNC (Thin Ethernet) and UTP as shown in Figure 65. You must select the Ethernet connector that you want to use with the switch that is at the back of the 8235.

Three Ethernet wiring schemes are supported:

- Thin (10Base2)
- Thick (10Base5)
- UTP (10Base-T)

When twisted-pair is selected, the LED next to the twisted-pair port on the rear panel of the 8235 Model 022 indicates the network status. Table 9 on page 74 summarizes what the various flashing patterns mean and what actions, if any, you should take.

Table 9. 8235 LED Error Code Flashing Patterns

LED Pattern	Meaning	Action to Take
On	Normal link is established.	None; normal operation.
Off	10Base-T is not selected.	Set the Ethernet connector switch to the 10Base-T (far left) position.
One flash	Link to 10Base-T is down.	Check that the hardware connections are secure. Re-establish the link.
Two flashes	Jabber error (possibly transient). The 10Base-T transceiver has detected a continuous frame transmission of 131 milliseconds or greater by the LAN controller in the 8235 Model 2. Transmission on the network is inhibited.	Wait a few seconds to see whether the problem goes away. If not, restart the 8235 Model 2, or contact IBM Product Support.

2.1.8.2 8235 Code Structure

The software that runs in the 8235 server can be separated into three pieces:

- Boot PROM
- Virtual ROM (VROM)
- The main software image

Boot PROM: The Boot PROM resides in ROM and performs the function of downloading a software image if there is no valid image in the VROM. Otherwise, the VROM performs software downloads. The Boot PROM accomplishes software downloads via Boot Protocol (BOOTP) and trivial file transfer protocol (TFTP) or via SPX. In addition to software downloads, the Boot PROM performs power-on-self test (POST) and switches the device to diagnostic mode if the POST fails.

VROM: The VROM serves to isolate the mainline programs from the hardware by providing the following:

- Device drivers for LAN and serial port I/O
- Buffer and memory management
- Management of non-volatile storage
- LED manipulation
- Message logging
- Acquiring VROM maintained data
- Acquiring hardware configuration information

The VROM also contains a bootstrap application that is capable of acquiring a new download by unattended BOOTP and TFTP or a NetWare SPX download from the Management Facility. The 8235 downloads new images through the LAN port (token-ring or Ethernet).

Main Software Image: The bulk of the run-time function in the 8235 is contained in the main software image. This image consists of the software kernel, frame forwarding support, management, and security.

2.1.8.3 Updating Microcode

The system structure for the 8235 makes it an excellent platform for future enhancements that can be obtained via software updates.

Downloading Modes: The 8235 can be put into several different boot up sequences under the control of one of the following:

- Management Facility
- Command shell
- Physical interruption (power on and off, pin reset)

The different modes are described in the following paragraphs.

Warm Boot: Under normal circumstances, the 8235 will contain a software image and configuration that has been stored in battery-backed RAM. When the system is rebooted (powered on or restarted due to a configuration change), it goes through a normal cycle. During this cycle, it will temporarily appear to the Management Facility to be in download mode. The device list window will indicate that the device is in DL mode. This condition should last for only a few seconds. If for some reason the 8235 has lost its code image or has been pin reset, it will remain in download mode until a management entity has loaded new code.

Download Code Only: The 8235 can be instructed to download a new code image only by issuing a Download command from the Management Facility. This means that it will load a new code image, but will maintain its configuration data.

Clear and Download: A Clear and Download command from the Management Facility will put the 8235 into download mode from the Boot Prom on the 8235 and will load both code and VROM, and will cause any configuration data in the 8235 to be lost. It will remain in download mode until a management entity loads a new version of code.

Pin Reset Switch: The 8235 has a tiny pinhole at the back that is not labeled. It is a pin reset that corresponds to an internal switch that performs the hard reset of the 8235 and is often overlooked. It should be used if you lose contact with the Management Facility due to hardware problems or if you lose the administrator's password. It performs the same function as the Clear and Download command. No indication of this pin reset is noted on the hardware itself.

2.1.9 Models Summary

The main difference between all the 8235 models is the communication port that is used.

Table 10 (Page 1 of 2). 8235 Models

Model Feature	Token-Ring	Ethernet	HS Serial Port (115.2 kbps)	Internal Modem	Serial Port (57.6 kbps)
8235-021	X		X		
8235-022		X	X		
8235-031	X		1-8	1-8	1-8

Table 10 (Page 2 of 2). 8235 Models					
Model Feature	**Token-Ring**	**Ethernet**	**HS Serial Port (115.2 kbps)**	**Internal Modem**	**Serial Port (57.6 kbps)**
8235-032		X	1-8	1-8	1-8
8250 module	X		X		
8250 module		X	X		

Note

Models 031 and 032 have empty slots into which you can install up to eight cards: eight modem cards, eight serial cards, or a combination of both.

2.1.10 Communication Options

Here is a brief description of the different communication options that the 8235 has:

- Models 021 (token-ring) and 022 (Ethernet)

 The new, high-speed base models, 021 and 022, support serial port speeds up to 115.2 kbps, enhancing the 8235 model offerings. These new models are shipped with eight RS-232-D (V.24/V.28) ports for attachment of up to eight modems with 115.2 kbps serial port speed. Excellent performance can be achieved with the high-speed V.34 data compression modems.

- Models 031 (token-ring) and 032 (Ethernet)

 These models do not contain a fixed port configuration. The customer configures the ports to meet their needs with any combination of modems and/or serial cards.

 Model 031 is an unpopulated token-ring base server, and Model 032 is an unpopulated Ethernet base server. Both models provide plug-in slots for V.34 modem cards and serial cards. These models support a total of eight cards (eight modem cards, eight serial cards, or a combination of both cards totaling eight).

 These models can support eight remote users simultaneously with reliable asynchronous transmission speeds up to 115.2 kbps. With the serial cards, you can configure some or all of the ports to attach external asynchronous terminal adapters for digital services, such as ISDN or Switched 56.

 The Management Facility of 8235 Models 031 and 032 is an extension to the facility provided with the other models of the 8235 and is enhanced to include management of the new V.34 integrated modems and serial cards.

 IBM has extended the flexibility of the IBM 8235 Models 031 and 032 remote access server with several new upgrade modules:

 IBM 8235-031 and 032 BRI module

 – 2B+D with V.110 and V.120 rate adaption.
 – S/T and U interface versions are available.
 – BRI module can be monitored from IBM MF. Configuration setup, revisions, and troubleshooting can all be managed remotely.

 IBM 8235-031 and 032 Sync/Async module

- User can connect synchronous devices (ISDN BRI TAs, CSU/DSUs and modem eliminators) directly to the IBM 8235/Models 031 and 032. The direct synchronous connection takes advantage of the faster line speed (128 kbps vs. 115 kbps), the elimination of extra timing bits (async has two extra timing bits per character transmitted), and the overhead of converting synchronous transmission into asynchronous transmission.
- Supports either synchronous or asynchronous communications channels.

• 8250 Modules

 These modules integrate IBM 8235 remote LAN access server product functions into the 8250 hub.

 There are two kinds of 8235 modules:

 - One for attaching an Ethernet network
 - One for token-ring network attachment

 These modules occupy a single slot in the 8250 hub chassis. The Ethernet module provides one Ethernet attachment switchable to any of the three Ethernet segments on the 8250 backplane. Likewise, the token-ring module provides one token-ring attachment that can operate at either 4 or 16 Mbps. The attachment is switchable to any of the seven token-ring backplane segments.

 Each module has eight serial communication ports. Each port has an RS-232-D (V.24/V.28) interface with a DIN connector for attachment to standard asynchronous modems. Data transfer speed ranges from 2400 bps up to 28.8 kbps, or even up to 115.2 kbps when using high-speed data compression modems. The modules come with eight DIN-to-25 pin RS232 patch cables to attach to external modems.

2.1.11 Supported Protocols

The 8235 supports remote clients using any of all the following protocols.

2.1.11.1 NetBIOS and 802.2

The 8235 software filters on LLC service access point (SAPs) and on NetBIOS names based on the filter tables contained in the server. The tables will be set up in the box, but the information can be overridden using the operating system shell. There are no external parameters available to manage filtering as there are for an IBM Token-Ring Bridge or for LAN Distance software. LLC SAP filters allow X′02, X′04, X′05, X′08, X′E0, X′F0 and X′F4 SAPs to be bridged. These are also configurable.

Frame forwarding (that is, the process of forwarding data from the client workstation to the LAN and from the LAN to the client) is accomplished differently depending on the protocol selected during the configuration of the connections.

2.1.11.2 Bridging

The token-ring acts like an IBM token-ring bridge with the NetBIOS and 802.2 protocols as shown in Figure 66 on page 78.

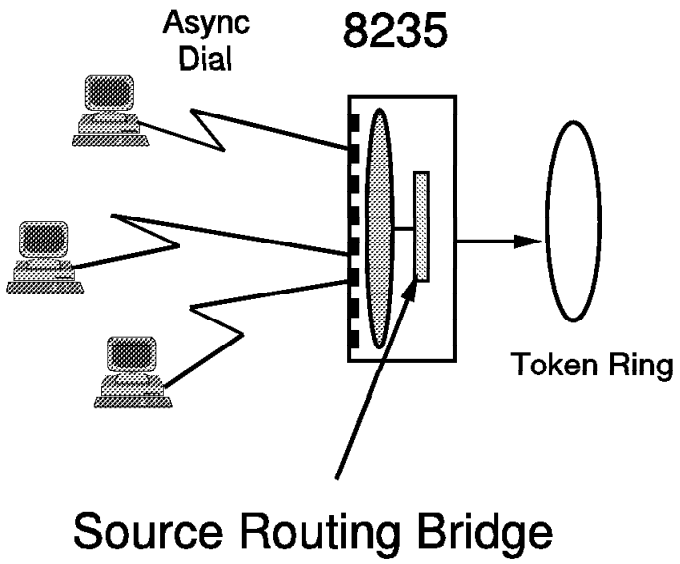


Figure 66. Source Routing Bridge

The bridged frames appear on the ring as if they came from an adapter. NetBIOS and 802.2 dial-in also supports specialized filtering to protect clients from broadcast traffic on the dial-in links.

The 8235 acts like a transparent bridge for Ethernet as shown in Figure 67.

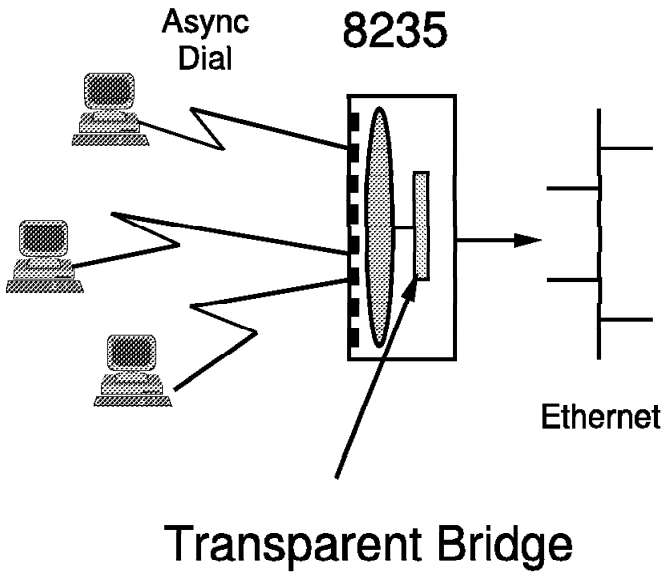


Figure 67. 8235 Acting As a Transparent Bridge

2.1.11.3 Ring Parameter Server

The ring parameter server (RPS) function has been implemented in the case where the 8235 is the only bridge on the ring. Here is an explanation of what the RPS function provides.

The RPS is the target for all request initialization MAC frames that are sent by ring stations during their attachment to the ring segment. The RPS function makes the following parameters available to all ring stations on the ring in response to the request initialization MAC frame:

- Ring number
- Ring station soft error report time value (default of 2 seconds)
- Physical location (not currently implemented)

There can be more than one RPS function active on any given ring segment.

Note

This differs from an IBM source routing bridge in that LAN reporting mechanism functions are not present in the 8235 which would allow it to report configuration information to LAN Network Manager (LNM) or to accept configuration changes from LNM.

2.1.11.4 IP Traffic

The 8235 will transparently forward IP traffic based on the IP address. The 8235 implements the proxy address resolution protocol (ARP) function to reduce broadcast traffic over the remote lines.

Note

This means that the 8235 will respond to all ARP queries for remote client addresses with its own hardware address instead of having the ARPs go across the WAN. The source stations will then forward packets from the remote clients to the 8235's physical address. The 8235 will then route the packet to the correct client based on the IP address.

An example of how the network would appear is shown in Figure 68 on page 80.

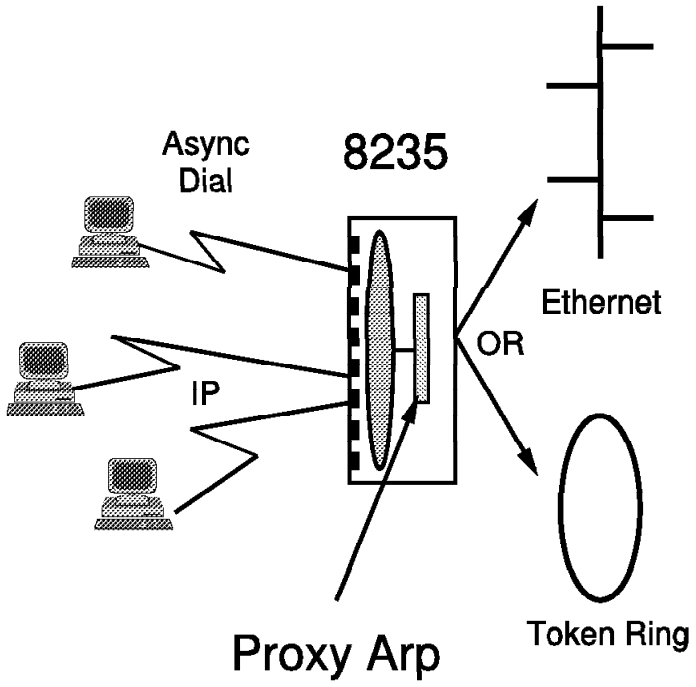


Figure 68. 8235 Proxy ARP

The 8235 will implement the following IP functions:

- IP Address Resolution Protocol (ARP)
- Internet Protocol (IP)
- Internet Control Message Protocol (ICMP)
- Transmission Control Protocol (TCP)
- User Datagram Protocol (UDP)
- Trivial File Transfer Protocol (TFTP)
- Boot Protocol (BOOTP)
- Telnet
- Routing Information Protocol (RIP)

For IP traffic, Van Jacobson Header compression is supported. This is transparent to the user, but enhances performance over the telephone network connection.

IP environments pose a unique challenge to dial-in access, as the addresses contain the identification of the network. If the users provide their own IP address, then they are limited to dialing in to the network for which they have been preconfigured. There are, however, some environments where the user will be dial in to the same network all of the time and want to keep the same IP address. Furthermore, because of the nature of IP address discovery (ARP), it is desirable to limit the amount of ARP traffic across the WAN.

Because of this, the 8235 supports address assignment in two ways:

1. Proxy ARP with static client addressing, which has the following properties:

- Dial-in client has configured IP address, provided to the box by IPCP.
- A user must dial-in or attach to the same network all of the time.
- Full end-user TCP/IP application suite support.
- IP address for each dial in client is resolved to MAC address of the LAN port (proxy ARP).
- Packets are routed based on host ID. If the network ID does not match the host ID, the packets will not be forwarded.
- Remote-to-remote is a special case. The 8235 recognizes it and forwards the traffic as a special case.
- Header compression is supported.

2. Proxy ARP with dynamic client addressing, which has the following properties:
 - The 8235 provides unique client IP address through IPCP.
 - Dial-in user can dial into any network that is reachable from the LAN to which the 8235 is connected.
 - The user does not own a well-known IP address. While this may prohibit the use of dial-in clients as servers, it allows the use of most user-oriented software.
 - IP address for each dial-in client is resolved to MAC address of LAN port.
 - Packets are routed based on host ID.
 - Remote-to-remote is a special case. The 8235 recognizes it and forwards the traffic as a special case.
 - Header compression is supported.

Note

The IP address of the 8235 box itself can only be assigned through the Management Facility.

2.1.11.5 IPX Traffic

The 8235 implements an IPX router function as defined by Novell (see Figure 69 on page 82).

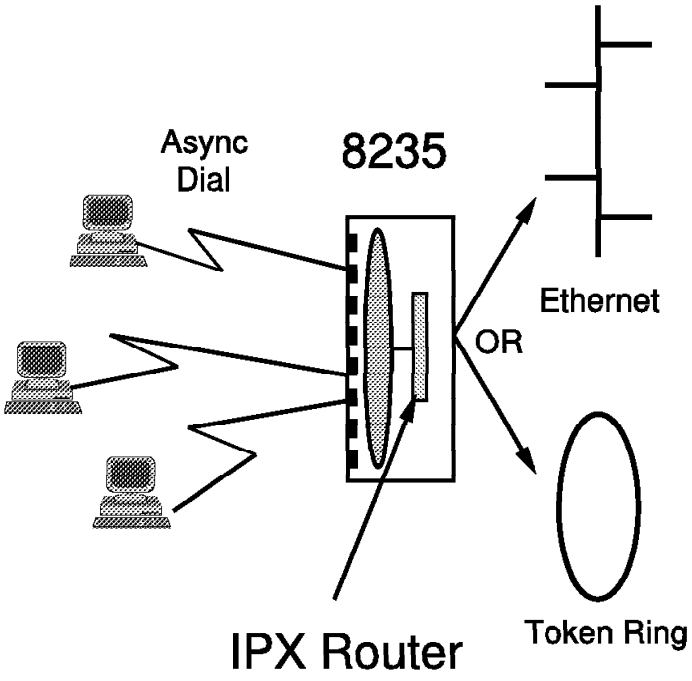


Figure 69. 8235 IPX Router

Basic IPX protocols implemented by the 8235 are:

- Internet packet exchange (IPX) providing the basic network layer transport for NetWare IPX.
- Sequenced Packet eXchange (SPX) for reliable byte stream protocol. This is used for NetWare diagnostics and for downloading code images over IPX.
- Routing information protocol (RIP) which provides a mechanism for IPX routers to exchange network topology information as needed to maintain routing tables. RIP uses a distance vector algorithm to calculate the best routes.
- Service advertising protocol (SAP), which provides a mechanism for end systems to locate NetWare services. The 8235 advertises its management via SAP.

The 8235 supports dial-in routing by the remote user for IPX onto the local LAN. The network number of the dial-in port can be assigned by the administrator. If the assigned number is in use on the network when a user dials in, the box can be configured to take one of three actions: use the net number anyway, use a random number, or refuse the connection. If the dial-in client uses a non-zero node address, the server will accept it. If the client uses a zero node address, the server will provide the client's address. The 8235 supports the following IPX frame types:

- Ethernet II (Ethernet)
- 802.3 (Ethernet)
- 802.2 (Ethernet)
- SNAP (Ethernet)

- SNAP (token-ring)
- 802.2 (token-ring)

2.1.11.6 AppleTalk ARA 2.0

You can configure the 8235 as an end node or router and assign it to an AppleTalk zone.

AppleTalk protocols support zones for managing user access to network devices and services. Zones are logical names associated with networks. The network administrator chooses an AppleTalk Phase 2 default zone during the initial setup of the network. The 8235 can be placed in this default zone or in a valid Phase 2 zone in the zone list.

Note: The 8235 supports AppleTalk Phase 2 networks only.

The 8235 may appear as one of the following on the AppleTalk network:

- A node
- A router

End Nodes: Apple Remote Access (ARA) software allows Apple users to connect to an AppleTalk network through a modem/serial link. The ARA remote client calls a locally attached ARA server. The ARA server provides the client with access to LAN resources (electronic mail, file servers, printers, and network applications).

An ARA server operating in end-node mode is responsible for forwarding packets sent to and from the ARA client. The ARA server examines packets sent on the network. If the destination is the ARA server or a remote ARA client, or it is a broadcast packet, then the server accepts the packet. If the destination is a remote ARA client, the server sends the packet across the serial link to the remote client.

AppleTalk remote access protocol (ARAP) requires the ARA server to prevent broadcast routing table maintenance protocol (RTMP) information from being forwarded to the client over the serial link. The ARA client does not need the RTMP broadcast information.

A packet sent from an ARA client to a user on a different network is forwarded by the ARA server to a router using the most recent router method. This method is used because the ARA server operating in end-node mode is not a router and must forward the packet based on the most recent information it has received about the destination. The most recent router method does not ensure the packet is routed to its destination by the fastest available path. The ARA server in end-node mode provides for easy configuration. An end node does not require a new (additional) network number and is less intrusive on large networks because it does not broadcast RTMP packets as a router does.

Advantages of using the 8235 in end-node mode

- Easy setup.
- Network number not required.
- Serial link traffic could be minimized:
 - NBP broadcasts not destined for the client are not forwarded.

– RTMP packets are not forwarded. The 8235 is not a router in this mode.

The end node implementation of ARAP in the 8235 is compatible with Apple′s ARAP implementation. When the 8235 is configured to function as an end node, the 8235 forwards the data packets to and from the ARA clients in the same way as an ARA server.

With the 8235 functioning as an end node, all 8235s on the network can be assigned to one zone in the Phase 2 zone list with the 8235 appears in option. Network administrators would only need to access one zone to find all of the 8235s on the network.

8235 ARA clients can be assigned to a different Phase 2 zone. Assigning ARA users to a different zone can help reduce NBP broadcasts over the serial link if the zone chosen does not receive many NBP broadcasts. This can significantly improve performance over the serial link.

ARA Routers: An ARA server in router mode acts as a router between two networks: the local internetwork on which the server resides and a network into which remote clients are assigned. In contrast to an ARA end-node server, which makes a remote ARA client a node on the network, an ARA server in router mode makes an ARA client a node on a separate dial-in (remote) network. The dial-in network has as many nodes as there are ARA clients connected to the server. This ARA client network can be assigned to any zone on the network including a zone in the Phase 2 zone a list or a newly created zone.

When acting as a router, the ARA server maintains complete zone and routing tables of the internetwork in memory. When a node on the internetwork sends a packet, the router examines the packet header and determines the destination by checking the routing table. If the destination is a remote ARA client, the packet is routed to the dial in network and sent to the node number of the ARA client.

When a packet is sent from an ARA client to the local network over the serial link, the ARA server uses its routing table information to route the packet to its destination by the most efficient path in the routing table.

An ARA server configured as a router can isolate the ARA client from AppleTalk broadcast packets by permitting the client to be located in a dial-in zone. This improves performance over the serial link, because only broadcasts into the dial-in zone are sent over the serial link.

Advantages Using the 8235 in Router Mode: The 8235 can be configured to function as a conforming router or as a seed router. A conforming router obtains routing information from other routers on the network. A seed router provides the routing information to the other routers on the network.

The 8235 operating in router mode provides some advantages:

- AppleTalk broadcast packets sent over the remote link can be limited by placing the remote link into a dial-in zone. Only broadcasts into that zone are sent over the link.
- The 8235 knows the fastest route to all networks and will route client packets by the most efficient path.

- The 8235 can be assigned to a different zone in the Phase 2 zone list. By assigning all 8235s to a particular management zone, network administrators only need to access one zone to find all 8235s on the network.
- The 8235 can isolate ARA clients from the rest of the Internet by assigning clients to a dial-in zone. Each client has a different node number in this zone. The dial-in zone may be a newly created zone. It does not have to be in the Phase 2 zone list. All dial-in clients can be placed into this dial-in zone. Network administrators can monitor dial-in activity by monitoring this zone.
- Network and zone information is configurable for ARA clients.
- For LAN-to-LAN connections, the 8235 must be in router mode.

IP Information: IP forwarding allows the 8235 to provide IP address assignments for dial-in clients. The client's IP address must be part of the Ethernet/IP network. Other IP hosts on the network communicate with the dial-in users through the 8235. The 8235 responds to Address Resolution Protocol (ARP) requests that are destined for a client IP address. This is referred to as proxy ARP. When an IP host requests an 8235 client IP address, the 8235 responds to the host with its own Ethernet address, specified on the IP configuration page. The 8235 accepts client packets and forwards the packet to the correct IP client/address.

IP packets are routed across an AppleTalk network by means of encapsulation. The 8235 sends IP packets to Macintosh dial-in clients by encapsulating the IP packet within an AppleTalk packet. The 8235 forwards IP packets from an ARA client to an IP host by de-encapsulating the IP packet.

The 8235 ARA dial-in clients appear as if they are directly connected nodes within the IP network. The IP host and the dial-in client are not affected by the fact that their packets are being routed through the 8235.

The Macintosh dial-in client uses the name binding protocol (NBP) to search for an IPGATEWAY device type in a specified zone. Since the 8235 is the ARA server for the client, the 8235 processes all of the client's AppleTalk packets and checks its configuration to see if it is configured as an IP gateway for that zone. If it is, the 8235 responds to the Macintosh dial-in client that it is an IPGATEWAY.

The dial-in client sends a kinetics internet protocol (KIP) command to the 8235 asking for an IP address. The 8235 responds with the dial-in client's IP address, subnet mask, broadcast address and the IP address of the name server.

To communicate with an IP host, the user must have an IP address. IP addresses are assigned to a Macintosh client as follows:

- Per user: When a dial-in connection is made, the 8235 checks the user list to see if there is a user IP address. If there is a user IP address in the user list, the 8235 assigns this IP address to the client.
- Per port: If there is no IP address in the user list, the 8235 assigns the port IP address to the client.

2.1.12 Security

The 8235 provides several security features. Passwords for both dial-in and LAN-to-LAN connections are automatically encrypted. User lists store user profiles which include user names, passwords, permissions and dial-back. If dial-back is selected in a user profile, the 8235 will hang up after the dial-in or LAN-to-LAN connection is established and then call the user back at a configured (fixed dial-back) number or at a number entered by the user when the connection was established (roaming dial-back). Unauthorized access to the 8235 device configuration or user list can be prevented by assigning the 8235 an administrator password. This password is stored in the 8235 device configuration information, not in the user list.

The 8235 has a unified security architecture which allows any security server on the LAN to be used to authenticate any user regardless of the protocol being used. This allows a centralized security method to be used for all authentications. 8235 Version 2.0 code or later supports three authentication databases:

- 8235 User List
- NetWare Bindery
- SecurID ACE/Server

The 8235 prompts separately for the user name and password for each method of authentication. Thus, more than one security method can be used simultaneously. SecurID could be used to authenticate an individual user who then logs into a NetWare Bindery group and is granted the access privileges associated with that group. Because the user protocol does not matter, the NetWare Bindery could be used to authenticate an Apple Remote Access (ARA) Version 2.0 dial-in user.

2.1.12.1 8235 User List

Using the 8235 Management Facility a user list can be created, edited, and then saved to a file or loaded into the 8235. The 8235 user list stores the names, passwords, and permissions of users authorized to dial into or out of the network or to connect to another network. User lists are stored in battery backed-up RAM in the 8235. Each 8235 can have a different user list or one user list can be downloaded to multiple 8235s. The NetWare Bindery or SecurID is recommended if there are more than 500 users.

2.1.12.2 Using the NetWare Bindery

The NetWare Bindery is a database that resides on a NetWare server. This database contains profiles of network users that define each user′s NetWare name, password, dial-back number, and the permissions to use one or more of the 8235 functions such as dial-in, dial-out or LAN-to-LAN.

When bindery authentication is enabled, it replaces the 8235 user list authentication.

With bindery security enabled, the bindery services utility can be used to create bindery groups for dial-in, dial-out, and LAN-to-LAN users. The group names are 8235_DIALIN, 8235_DIALOUT, and 8235_LAN-to-LAN. The bindery dial-in user groups are used when a user dials into the network using a NetWare name and password. The 8235 logs in to the NetWare server with this user name and

password and then logs out. If the 8235 logon to the server was successful, the 8235 allows the user to access the network through the 8235.

2.1.12.3 Bindery and Apple Remote Access (ARA)

To use the bindery, ARA Version 2.0 users must have the 8235 Security Module in their Macintosh system Extensions folder. This module supplies a security drop-in, which provides 8235 password encryption (thereby allowing bindery security to work with ARA Version 2.0).

2.1.12.4 Using SecurID

Security Dynamics, Inc. manufactures two security solutions that are compatible with the 8235. The first is a multiport, stand-alone device that can be inserted between the 8235 and the modem. This solution requires no particular configuration of the 8235. The device dialing in must be capable of handling the authentication dialog.

Macintosh users who have the external SecurID client box installed for their 8235 can still use their command control languages (CCL) as before; however, SecurID should not be enabled in the 8235 Management Facility, as this will trigger the 8235 internal SecurID client.

SDI′s second security solution is the Security Dynamics ACE/Server, which is a system of server and client software and SecurID cards. Once enabled, SecurID authentication is used for all protocols (IP, IPX, NetBEUI, 802.2 LLC, and ARA).

The 8235 can use SecurID to protect its serial ports from unauthorized dial-in access. SecurID authenticates users and may be used in conjunction with the 8235 user list or the NetWare Bindery. See Figure 70 for the SecurID configuration.

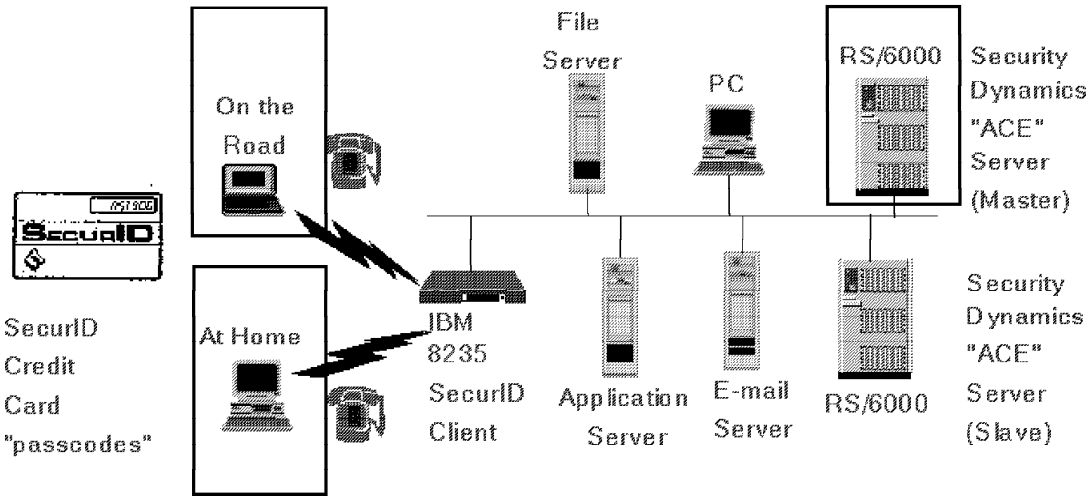


Figure 70. 8235 Security System

SecurID authentication is not required of dial-out users, users managing the 8235 with the command shell, or users managing the 8235 with the 8235 Management Facility. SecurID does not protect the 8235 from dial-out, LAN-to-LAN, or local

area network shell access. If the 8235 is using SecurID authentication, incoming LAN-to-LAN connections are not permitted.

The components of a full implementation of SecurID are as follows:

- SecurID server software

 This software runs on a UNIX machine. User data protocol (UDP) is used to communicate with the client software running on the 8235. This server software is purchased from Security Dynamics, Inc.

- SecurID client

 This is the component running on the 8235 that communicates with the SecurID server via UDP. It is compatible with SecurID server software Version 1.1 or later.

- SecurID card

 This component is a card that provides the user with a passcode number needed to access the SecurID server.

- Dial-in client software

 This is the standard 8235 Remote Dial-in Client Version 2.0 or later for PC users or Apple Remote Access (ARA) Client Version 2.0 or later for Macintosh users.

2.1.13 The Activity Logger

The activity logger runs under Microsoft Windows and DOS. It provides information about 8235s and their dial-in activity on the network.

The logger carries out the following tasks:

- It records the dial-in activity of the 8235 on the network.
- It notifies the network administrator of 8235 activity according to a set of priorities and classes selected by the administrator.

The 8235 logs its activity to another station using a mechanism of SNMP called a trap. Each time the 8235 logs an event, it sends a trap message to its trap host.

The trap host can be one of the following:

- A workstation running the 8235 Activity Logger
- An IP host with an SNMP manager

There can only be one trap host associated with an 8235 at any given time. This trap host is configured in the 8235 Management Facility on the SNMP configuration window. There are two host types to choose from: None and IP.

If you select IP, then you can also specify the IP address of the trap host. This IP host must be an SNMP manager and have some facility for displaying SNMP trap messages if it is to be used as the activity logger. For example, this could be a NetView for AIX management station.

If you select None, then the trap host address cannot be specified via the 8235 Management Facility. Instead, once the 8235 activity logger (which runs on top of IPX) selects an 8235 as a device to be logged to that workstation, the selected 8235 sends all of its trap messages to that workstation. If an 8235 is selected on one activity logger workstation while another Activity Logger workstation is the

current trap host, the new workstation becomes the new trap host. This provides flexibility in case a trap host goes down because it is easy to switch over to a backup host.

2.2 IBM 2210 Nways Multiprotocol Router

This section provides an overview to the IBM 2210, including a description of the hardware and an overview of the software package. Further information is found in the *IBM 2210 Nways Multiprotocol Router Maintenance Information*, SY27-0345 and the *IBM 2210 Nways Multiprotocol Router Planning and Setup Guide*, GA27-4068.

2.2.1 Models of the IBM 2210

The IBM 2210 is available in several models, based on the types of networks you want to support.

IBM withdrew the Models 121, 122, 123, 124, 125 and 126. Models 121, 122, 123 and 124 had one LAN port, two serial connections, 2 MB Flash and 4 MB DRAM and were replaced with the Models 12T and 12E. Models 125 and 126 had one LAN port, two serial connections, 2 MB Flash and 4 MB DRAM and were replaced with the Models 127 and 128.

Table 11 on page 90 shows the different models and the offerings of the IBM Nways Multiprotocol Routing Network Services that are available.

The only differences between some of the models is the amount of flash memory and DRAM. Flash memory is used to store a compressed version of the router′s software while DRAM memory provides the working memory for the router programs and the router network tables.

Note: Flash memory is not able to be upgraded on the 12x models of the IBM 2210.

You can add an additional 4 MB of flash memory to the 14T and 24x models of the IBM 2210 by replacing the installed flash memory with an 8 MB Memory Expansion Feature. This upgrade provides a total of 8 MB of flash memory for those models.

If you want to maintain multiple copies of software for various releases, you may want to consider a model with 4 MB of flash memory.

IBM 2210′s DRAM provides the working memory for the router programs and the router network tables. The amount of required DRAM in an IBM 2210 is determined by the size and complexity of the network that the IBM 2210 must support.

You can upgrade the DRAM on all models of the IBM 2210 to a maximum of 16 MB using IBM′s 16 MB Memory Expansion Feature.

Certain models of the IBM 2210 support ISDN. You cannot use one of the standard WAN ports for ISDN. Software support for ISDN must be ordered separately.

Table 11. IBM 2210 Models

Hardware						Software		
Model	**LAN**	**No. of WANs (see note)**	**Flash Memory**	**DRAM**	**ISDN**	**Base**	**Additional Routing**	**ISDN**
12T	Token-Ring	2	4 MB	4 MB		x	x	
12E	Ethernet	2	4 MB	4 MB		x	x	
127	Token-Ring	2	4 MB	4 MB	x	x	x	x
128	Ethernet	2	4 MB	4 MB	x	x	x	x
14T	Token-Ring	4	4 MB	8 MB	x	x	x	x
24T	2 Token-Ring	4	4 MB	8 MB	x	x	x	x
24E	2 Ethernet	4	4 MB	8 MB	x	x	x	x
24M	1 Token-Ring 1 Ethernet	4	4 MB	8 MB	x	x	x	x

Note: The standard WAN ports on the IBM 2210 will support any of these physical interfaces:

- EIA RS 232-D/V.24
- V.35
- V.36
- X.21

The ports of the different models are shown from Figure 71 through Figure 74 on page 91. The models shown in each figure differ only in the amount of DRAM and flash memory they contain, as described above.

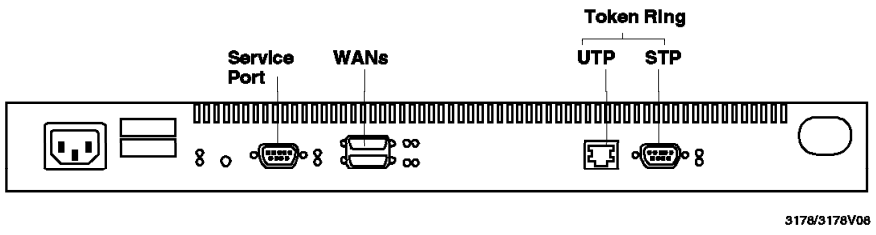


Figure 71. Model 12T

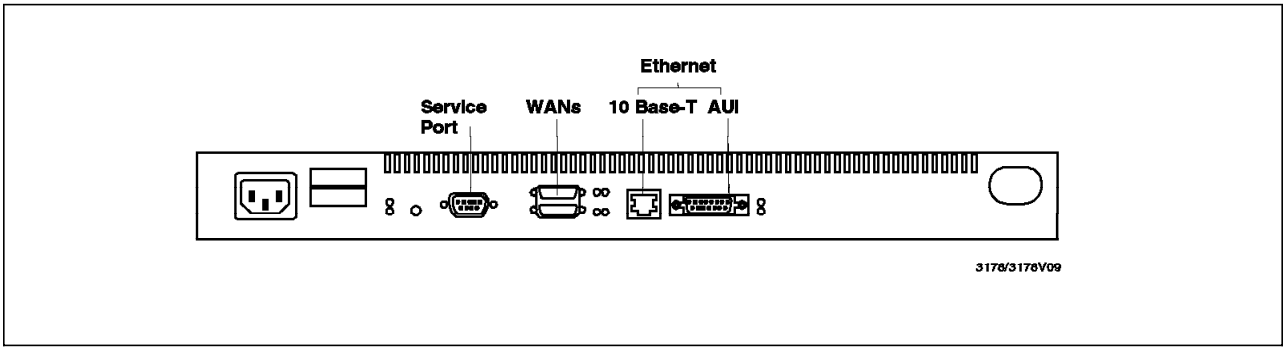


Figure 72. Model 12E

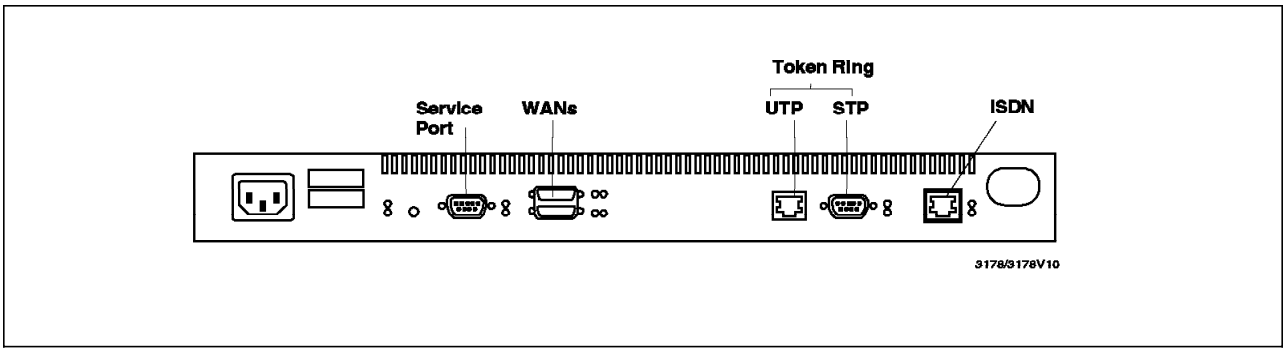


Figure 73. Model 127

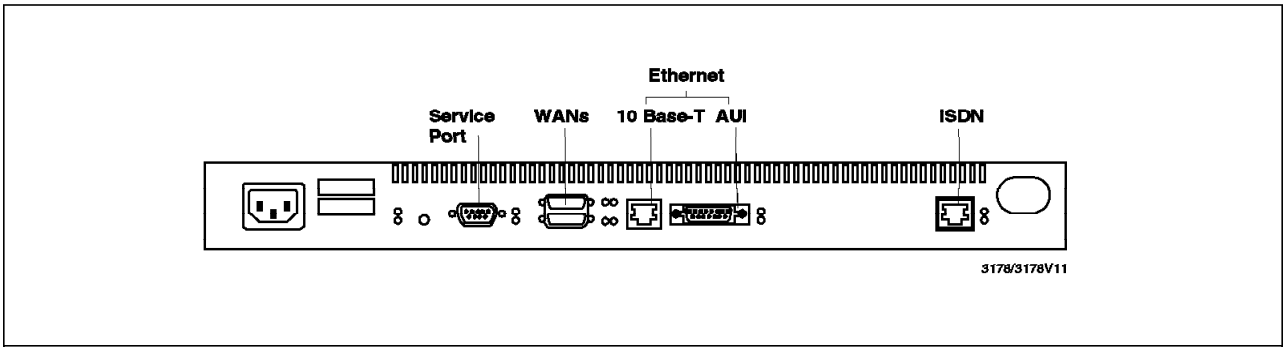


Figure 74. Model 128

2.2.2 Indicators on the IBM 2210

The IBM 2210 has green and amber LEDs that indicate the status of the system and of individual ports. Green indicates normal operation; amber indicates a problem.

The LEDs are on both the front and the back of the IBM 2210, so you can place it with either side facing forward. This is shown in Figure 75 on page 92 and Figure 76 on page 92.

Note: The figures shown are for Model 12T. The port LEDs are specific to each model.

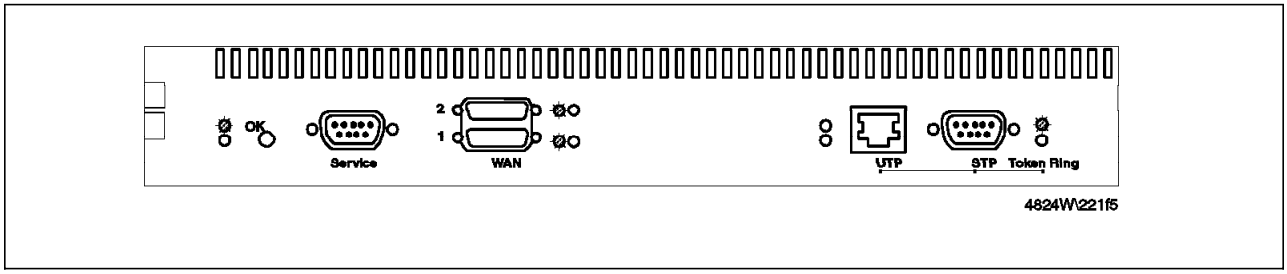


Figure 75. LEDs on the Port Side of Model 12T

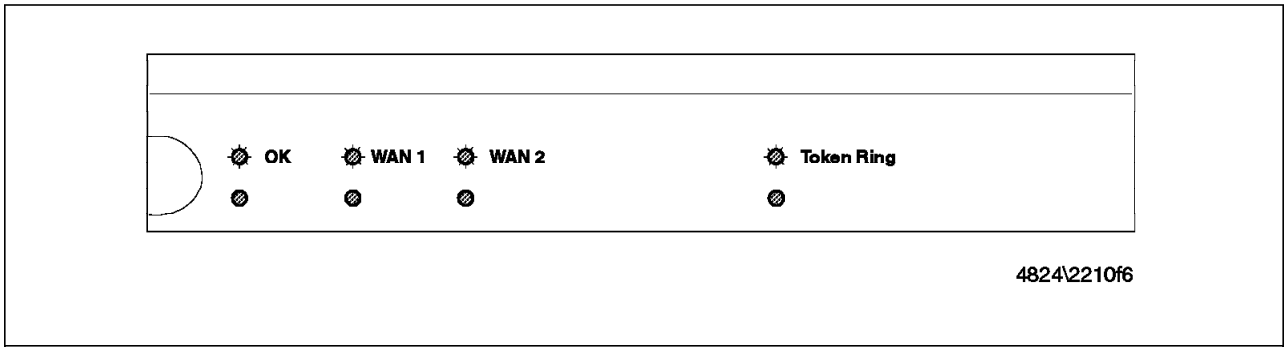


Figure 76. LEDs on the Side Opposite the Ports of Model 12T

2.2.3 The Reset Button on the IBM 2210

If you press the reset button, it will reload the operational code. Also, if you press this button within 10 seconds of powering on, the 2210 will enter the extended power-on self-test (POST). Extended POST allows you to test the memory more extensively than POST.

The reset button on the IBM 2210 is recessed to prevent accidental activation and is shown in Figure 77.

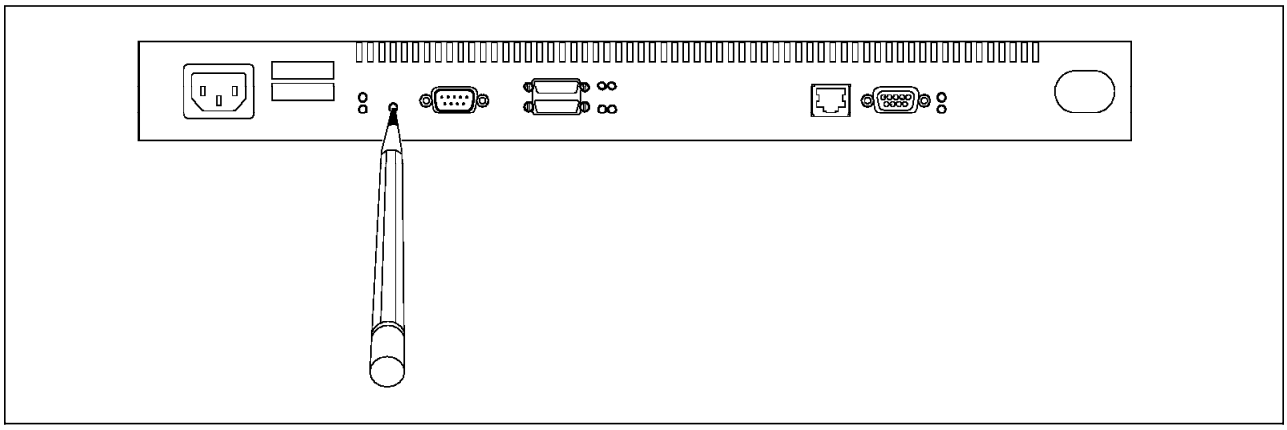


Figure 77. Reset Button on the IBM 2210

2.2.4 Networks Supported by the IBM 2210

The IBM 2210 supports the following LAN connections:

- Token-Ring (IEEE 802.5) with STP or UTP connection
- Ethernet (IEEE 802.3) with AUI or 10Base-T connection

Every IBM 2210 supports the following serial connections:

- EIA 232D/V.24
- V.35
- V.36
- X.21

Note: RS449 is also supported using the V.36 cable available for the IBM 2210.

In addition to these serial connections, you can order optional support for ISDN.

2.2.5 Accessing the IBM 2210

You can access the IBM 2210 using the following methods:

- An ASCII terminal (or emulator) attached directly to the service port
- An ASCII terminal (or emulator) attached via a modem to the service port
- A Telnet session

2.2.5.1 Local Access

You can access the IBM 2210 locally through its service port, using an ASCII terminal or emulator. The DEC VT100 terminal is supported, as well as devices that are configured to emulate it. The settings should be:

- No parity
- 8-bit word length
- 1 stop bit
- 300 bps-38.4 kbps bit rate

The IBM 3101, 3151 and 3161 display stations are also supported. For further information on these, please refer to The IBM 2210 Nways *Multiprotocol Router Planning and Setup Guide, GA27-4068*.

2.2.5.2 Remote Access

You can access the IBM 2210 remotely using either Telnet or a terminal attached to the service port via a modem.

The modem must use asynchronous operation and support the AT command set. The modem connected to the IBM 2210 must be set to auto-answer mode.

2.2.6 Software Package

Nways Multiprotocol Routing Network Services (MRNS) is the software that runs on the IBM 2210 and it comes as a base package, plus two separately orderable packages - one containing support for additional routing protocols and the other containing the ISDN support. The protocols supported by each package are:

- Base offering

 - TCP/IP over point-to-point (PPP), frame relay, and X.25
 - Bridging over PPP
 - Source-routing bridge (SRB)
 - Transparent bridge (TB)
 - Source-routing transparent bridge (SRT)
 - Source-routing - translational bridge (SR-TB)
 - SNA/DLSw over PPP, frame relay, X.25, and SDLC
 - Bandwidth reservation for PPP
- Additional Routing Protocols Feature
 - Internetwork Packet Exchange (IPX) over PPP, frame relay, and X.25
 - AppleTalk over PPP
- ISDN Feature
 - Supported over IP, IPX, AppleTalk, SRB, TB, SRT, and SNA/DLS

2.2.7 MRNS Overview

This section provides an overview of the Nways Multiprotocol Routing Network Services (MRNS) software for the IBM 2210. It includes descriptions of the boot process, the user interface and the event logging system (ELS). Further information can be found in the *Nways MRNS Software User's Guide*.

The Nways MRNS is the software that supports the IBM 2210. The Nways MRNS has three components:

- The code that provides the routing, bridging, data link switching, and SNMP agent functions for the IBM 2210
- The configuration program, which offers a graphical user interface that allows you to configure the IBM 2210 from a workstation
- A monitoring system that allows you to perform network management, problem determination, and configuration

2.2.7.1 Boot Files and Boot Processes

The IBM 2210 does not have a hard drive like the 6611 Network Processor, so it needs another method to load its operating system (referred to here as the boot file).

The boot file can be loaded (booted) from the following sources:

1. Flash memory referred to as the integrated boot device (IBD).
2. An external server which supports the TFTP server function. This could be another router which supports the TFTP server function (such as another IBM 2210).
3. The console port using ZModem.

Note: The IBM 2210 is delivered preloaded with a boot file in the IBD.

The IBM 2210 has a boot configuration database which holds information on all available boot files. Each entry in the database contains the location of the server host where the boot file resides and the path, file name, and a timeout

value for the boot file. You can add entries to the database by using the following command:

```
Boot Config>add boot-entries
```

On startup, the IBM 2210 will normally load itself with the boot file stored in the IBD, but it can use the boot configuration database to obtain a copy from a TFTP server should this boot file become corrupted or unusable.

The IBM 2210 may also use the boot protocol (BOOTP) to obtain its boot file, and uses the BOOTP client function to do so. The IBM 2210 will use the BOOTP protocol to learn its own IP address and the location (TFTP server) from which the boot file is obtained. It will then use TFTP to load the boot file from the TFTP server.

In order to cause the IBM 2210 to act as a BOOTP client, the interfaces over which the BOOTP packet should be broadcast are indicated by using the following command:

```
Boot Config>add bp-device
```

Note

When the IBM 2210 obtains its boot file at boot time from an external source, it loads the boot file into executable memory. However, it does not save
a copy in the IBD. If you want to move a copy into the IBD, you need to issue the following commands:

```
Boot Config>Copy Config or
Boot Config>TFTP get
```

Both commands use the TFTP protocol. The only difference is the format in which you specify the location of the file to be transferred.

The IBM 2210 does not allow you to initiate a transfer from another device to the IBM 2210, so you will need to start the transfer from the router operator's console.

The ZModem boot allows you to load router code through the console port using an ASCII terminal emulator package that supports the ZModem protocol. To load the code via this method, you enter:

```
>zb
```

The > prompt is the Boot prompt which is accessed by pressing Ctrl+C while the IBM 2210 is reloading.

Your ZModem software documentation will explain the commands required to start the upload.

2.2.7.2 IBM 2210 Configuration

The configuration process customizes the IBM 2210 for the network in which you intend to run it and the physical equipment being used. The configuration file may be created via the Nways MRNS Configuration Program and then transferred to the IBM 2210 or via commands entered at the operator console.

The configuration data resides in IBM 2210's non-volatile RAM (NVRAM) and is combined with the boot file when the IBM 2210 is restarted or reloaded, creating the operating environment of the IBM 2210. NVRAM is the only place from which the IBM 2210 will obtain the configuration information during a restart or reload.

Reloading the IBM 2210 causes the router to reload the boot file into RAM. At the same time, it customizes the operating environment using the configuration file on NVRAM.

To reload the IBM 2210, you issue the `Reload` command from the OPCON prompt.

Restarting the IBM 2210 doesn't cause the router to reload the boot file. It simply takes the configuration file on NVRAM and feeds it into the operating environment.

To restart the IBM 2210 you issue the `Restart` command from the OPCON prompt.

Changes made from the operator console configuration process (CONFIG) are immediately saved in NVRAM and, in most cases, will take effect once the IBM 2210 is restarted or reloaded. However, there are a few changes which will take effect immediately without the need to restart or reload.

Changes made from the operator console monitoring process (GWCON) take effect immediately. However, once the router is restarted or reloaded, these changes are lost. This facility could be useful if you wish to test some changes prior to making them permanent.

Note: The parameters which are changed from the GWCON process are a subset of the parameters which can be changed from the CONFIG process.

The Nways MRNS Configuration Program may also be used to configure the IBM 2210. The Nways MRNS Configuration Program runs under AIX, OS/2 and Windows and uses a GUI interface. When configuring via the Nways MRNS Configuration Program, you create a configuration file on the workstation which can be saved in two formats:

- An archive format which is stored in the workstation configuration database, and is readable by the Nways MRNS Configuration Program
- A 2210-readable format for transferring to the IBM 2210 via TFTP

Note: The 2210-readable format cannot be reloaded into the Nways MRNS Configuration Program, so it is highly recommended that you save an archive copy before creating and sending a 2210-readable file to the router. The 2210-readable file must be manually transferred to the IBM 2210 using one of the following commands:

- `Boot Config>Copy Config`
- `Boot Config>TFTP get`
- `>zc`

If you choose to create your configuration on the IBM 2210 console, then you should save a copy of it on an external server in case the NVRAM fails or the file is corrupted. You do this with the following commands:

- `Boot Config>Copy Config`
- `Boot Config>TFTP put`

The >zc command allows you to load a configuration file via the console port using an ASCII terminal emulator that supports the ZModem protocol .

To access the > prompt, you need to press Ctrl+C while the router is reloading.

Your ZModem software documentation will explain the commands required to start the upload.

2.2.7.3 MRNS User Interface

Access the Nways MRNS user interface through an ASCII console or emulator, as mentioned in ″Accessing the IBM 2210″ in 2.6.5.

By default, when you connect to the IBM 2210 you will not be required to enter a user ID or password, and you will have access to all router functions and commands. However, for security reasons you may want the users to enter a user ID and password when they connect to the router.

2.2.7.4 The Event Logging System (ELS)

ELS is a monitoring system that manages messages logged as a result of router activity. Using ELS commands, you can configure the system such that you only see the messages you need. ELS uses the concepts of subsystem, event number, message text, logging level, and group to help you manage the messages you see.

Subsystem is a predefined name for a router component, such as an interface or protocol. For example, IP is the subsystem name for the IP protocol, and TKR is the subsystem name for the token-ring interface.

The ELS Config process is accessed by issuing the `Config>event` command.

You can obtain a complete list of the subsystem names by issuing the `ELS Config>list subsystem` command. The output shows the subsystem name, the number of events for the subsystem, and a description of the subsystem.

Event number is a predefined number assigned to each message within a subsystem. You can obtain a complete list of events for a particular subsystem by issuing the `ELS Config>list subsystem subsys` command, where subsys is the name of the particular subsystem in which you are interested.

For example, `ELS Config>list udp` will list all possible events in the UDP subsystem. The output shows the event number, the logging level and the message text.

The message text is the actual text related to the event that has occurred and is used along with the subsystem and event number when the message is displayed by the MONITOR process. The logging level is a predefined category to which each event will belong, and which indicates the importance of the event. Note, whenever you use the ELS Config>list subsystem subsys

command to list all of the events within a subsystem, the logging level for each event is displayed.

Group is a user-defined collection of events that is given a name. A group can consist of events from different subsystems and of different logging levels. Once you have created a group, you can use the group name to manipulate the events in the group as a whole.

The Nways MRNS Event Logging System Messages Guide also contains a complete list of all events for all subsystems and includes the logging level for each event.

2.2.7.5 The IBM 2210 Configuration Program

The IBM Nways Multiprotocol Routing Network Services Configuration Program allows you to perform a complete configuration of an IBM 2210 Nways Multiprotocol Router. The Configuration Program is run on a workstation and has a graphical user interface.

Before using the Configuration Program you must perform an initial configuration on the 2210 to allow you to transfer these settings across to the IBM 2210 Router. The minimum requirement is that IP Routing is enabled to use the Trivial File Transfer Protocol (TFTP) or IP and SNMP are enabled to use the Communication option within the configuration program.

An Overview of the IBM 2210 Configuration Program: The IBM 2210 Configuration Program consists of two main windows:

- The Navigation window
- The Configuration window

The Navigation window displays a directory tree, consisting of the various components that you can configure.

To select any particular configuration screen, click the left mouse button on the item in which you are interested. The Configuration window will now display the configuration screen you have selected.

Help is available for each field within a panel. You may access the help by pressing PF1.

If the field requires you to enter a value, be sure you press CR (Enter/Return) after entering your value. If you don′t do this, the value may not be saved.

Hardware and Software Requirements: The following hardware is required to run the Configuration Program on the RISC System/6000 workstation:

- IBM AIX 3.1.5 or higher with Transmission Control Protocol/Internet Protocol (TCP/IP) enabled

 Note: AIX 4.0 and higher is not supported.
- IBM AIX windows
- 16 MB of memory
- A 3.5-inch diskette drive that can read and write
- 1.44 MB formatted diskettes
- 10 MB of available space on the fixed disk drive

- A graphics display that supports 640x480 resolution and 16 colors or gray scales
- A mouse

The following hardware and software are required to use the Configuration Program on a PS/2 workstation using an Intel 80386 or higher processor or a compatible system that has an Intel 80386 or higher processor.

For workstations running the Microsoft Windows program you need:

- IBM DOS 3.3 or higher, MS-DOS 3.3 or higher
- Microsoft Windows 3.1 or later versions
 - Win32s, included with the MRNS Configuration Program diskettes
 - WinSock 2.0 DLL (included with Win32s)
- TCP/IP application that uses WinSock 2.0 (this is only required for using the Configuration SEND function)
- 8 MB of memory
- 3.5 inch diskette drive that can read and write 1.44 MB formatted diskettes
- 10 MB of available space on the fixed disk drive
- A graphics display that supports 640x480 resolution and 16 colors or gray scales
- A mouse

For workstations running the IBM Operating System/2 (OS/2) Program, you need:

- OS/2 2.1 or later, including Warp
- IBM TCP/IP 1.2.1 or OS/2 or later (this is only required for using the Configuration SEND function)
- 10 MB of memory
- 3.5 inch diskette drive that can read and write 1.44 MB formatted diskettes
- 10 MB of available space on the fixed disk drive
- 10 MB of available swapper disk space on the swapper fixed disk drive partition
- A graphics display that supports 640x480 resolution and 16 colors or gray scales.

 Note: There is a known problem when running the Configuration Program on Warp. A selection of 65535 colors will prevent the program logo from displaying.
- A mouse

Anonymous FTP Site for the IBM 2210: IBM has established an anonymous FTP site for providing information and configuration program updates (and in the future other program updates) relating to the 2210.

The host name for the anonymous FTP site is nways.raleigh.ibm.com. If you have trouble resolving this name, the IP address is 192.35.236.5. After connecting to the machine, specify anonymous as the user ID and your e-mail address as your password. Check the README file on the anonymous FTP site in the /pub directory for the latest information.

The subdirectories where the Configuration Programs reside are as follows:

- /pub/config/2210/1.2.0.0/GA/diskettes for the diskette images
- /pub/config/2210/1.2.0.0/GA/runtime for the RISC System/6000 files

2.2.7.6 IBM Nways Multiprotocol Routing Network Services Release 3 - Enhancements

The MRNS Configuration Program Release 3 supports configuration for all of the functional enhancements for Nways Multiprotocol Routing Network Services Releases 1 and 2 and, in addition, offers the following:

- Support for the new 2210 Models 14T, 24T, 24E and 24M

 There are many packages of the MRNS Release 3 to support these new 2210 models or those currently available.

- Local LAN-to-LAN bridging support

 With the addition of multiple LAN connectivity on the new models, there is the obvious need for local bridging support. Users may configure LAN-to-LAN and LAN-to-WAN bridging using any of the following as appropriate:

 – Transparent bridging (TB)

 – Source-route bridging (SRB)

 – Source-route transparent bridging (SRT)

 – Source-route - Translational bridging (SR-TB)

- AIW Version 1 DLSw for SNA, and NOW NetBIOS, support (RFC 1795 compliant)

 MRNS′s DLSw is now compliant with RFC 1795, referred to as the AIW Version 1 DLSw. MRNS′s DLSw will still interoperate with the DLSw implementation in MRNS V1 R1 and R2 for SNA traffic but not for NetBIOS (prior releases support NetBIOS only via bridging).

- EasyStart, automatic configuration capability

 The goal of EasyStart is to eliminate the need for local initial configuration, essentially creating a ″plug and play″ installation.

 EasyStart allows network download of initial router configuration. When the system starts, and there is no configuration information, EasyStart attempts to obtain it from a network server. If EasyStart fails, the fall back is to use the local ASCII console.

 Once the initial configuration is retrieved from the network, the system is automatically restarted to cause the new configuration parameters to take effect.

- Data Compression over Point-to-Point Protocol (PPP)

 Support has been added for the draft standard PPP Compression Control Protocol and, currently, for a single data compression engine:

 – *Deflate - LZ77*

 PPP data compression is negotiated by PPP at link open time; the algorithm(s) used and the preference order can be set on pre-interface basis (once additional algorithms are introduced), to allow for control of the (substantial) memory usage of compression dictionaries (about 80 KB per direction with Deflate, 24 KB per interface with Stacker, over 90 KB per direction with BSD, and 64 KB per direction with Predictor).

PPP data compression can be used over any supported PPP interface, and can be used at the same time as Bandwidth Reservation (BRS will operate on data before compression is applied). When compression is in use, all data that passes over the interface is compressed. The impact of attempting to compress already compressed traffic varies according to the algorithm in use.

The compression achievable varies greatly according to the traffic. Using the Calgary Corpus standard of binaries, text files and image files, the Deflate algorithm achieves a ratio of 2.08:1. This compares to the following other algorithms:

- *Stacker-LZS: 1.82:1*
- *BSD Compress-LZW: 2.235:1*
- *Predictor: 1.67:1*

- LAN Network Manager (LNM) support

 The 2210 / MRNS LNM support is a source-route (SR) bridging option that enables LAN Network Manager agents on the 2210 bridge. The LNM function supports the following LNM agents:

 - *Configuration Report Server (CRS)*

 The CRS agent collects and reports MAC ring topology changes to the IBM LNM application. It will send out CRS MAC requests to query the status of other ring stations when requested by the LAN Network Manager.

 - *Ring Error Monitor (REM)*

 The REM agent collects MAC error reports from ring stations. When thresholds are exceeded, REM forwards error information to the LAN Network Manager.

 - *Ring Parameter Server (RPS)*

 The RPS agent services MAC requests from ring stations for ring parameter information and informs the LAN Network Manager of ring insertions.

- National ISDN-1, AT&T #5 ESS and Nortel's DMS-100 (US and Canada) supported on the 2210 ISDN Models 127 and 128

 The North American ISDN support is provided in Release 3 on the 2210 ISDN Models 127 and 128. With this support, users can attach the 2210 ISDN BRI port to one of the following:

 - AT&T #5 ESS switch
 - Nortel's DMS 100 switch

- WAN Re-Route

 The WAN Re-Route function is an enhancement to the IBM 2210 Multiprotocol Routing Network Services (MRNS) software. It allows the activation of an alternate network interface when a primary interface fails. WAN Re-Route is more flexible than the standard WAN Restoral feature (WRS) currently provided because the alternate link may have a different termination point than the primary link. It uses the dynamic routing abilities of the different routing protocols (IP RIP, IP OSPF, IPX RIP, etc.) or bridging protocols to find alternate paths through the new network topology. It also allows the backup

of all DLC types, that is frame relay, PPP and X.25, whereas WRS supports PPP links only.

- SNMP Enhancements

 As new functions are added to the MRNS, additional SNMP support is also necessary to ensure comprehensive network management capability. With Release 3, expanded SNMP MIB support has been added for SDLC links, LLC, BRS and the enhanced DLSw functions.

- IBM MRNS Configuration Program - Release 3 Enhancements

 The Release 3 MRNS Configuration Program enhancements include the following changes of the Release 2 support:

 – InARP support for IP, IPX and AppleTalk.

 – Ability to retrieve a configuration file from a 2210 and display its parameters.

 – Ability to create an ASCII flat file for printing purposes. The ability to import an ASCII file, verify contents and subsequently send to a 2210 is not yet available.

 – Drag and drop of certain lists.

 – Enhanced validation of file parameters.

- Additions to the Additional Routing Protocol Package

 – DECnet IV over PPP, frame relay (FR) and X.25 (2210 to 2210)

 Release 3 includes support over PPP data links as well as MRNS Release 2.

 – DECnet V / OSI protocols over PPP, FR and X.25

 The Digital Network Architecture (DNA) Phase V packet forwarder provides packet forwarding for 2210 routers in accordance with the Phase IV and Phase V router specifications of the DECnet protocol family. This allows a router to connect to systems using DECnet software (DNA Phase IV and Phase V network protocols) on different physical networks.

 – Banyan VINES over PPP

 Support of BVCP (Banyan VINES Control Protocol) over frame relay and X.25 (2210 to 2210) was initially offered in MRNS Release 2 and continues with Release 3. With MRNS Release 3, support of PPP data links is also provided. Because PPP is a nonproprietary protocol, the BVCP addition allows 2210 routers to interoperate with other vendor routers which abide by RFC 1763. Another advantage of the BVCP implementation is that we can expand VINES supports to any media that supports PPP.

- Optional Switch for Filtering Nonbridged Packets (Inbound Only)

 The switch is stored in SRAM and new user interface commands have been added to allow the customer to specify whether or not the nonbridged packers are filtered.

 A MAC filtering/bridging switch for nonbridged packets has been inserted, which allows the user to select whether nonbridged packets are filtered or not.

 The filtering of non-bridged packets will only occur when the following conditions are met:

 – Bridging is enabled.

– For inbound packets only (that is, packets coming from a LAN segment and not from a WAN interface).

– When the switch is set to allow filtering of non-bridged packets and when the filter parameter indicates the packet should be filtered.

MRNS, together with the IBM 2210 Nways Multiprotocol Router, provides users with a broad range of networking products and services for high-speed, integrated, manageable, and open networks. The 2210 Nways Multiprotocol Router connects local-area and wide-area networks to form a physically integrated network that transports multiple networking protocols between applications speaking the same protocol.

Plans are in process to eliminate sending copies of backup media diskettes since the desired software package is preloaded on the 2210. Instead, current licenses provide instructions on how to retrieve a copy of the code via Internet access to the MRNS Code Server.

Note:

AIW is the APPN Implementers Workshop who support the DLSw Related Interest Group (RIG) that evolved the RFC 1795 standard.

DECnet IV over FR and X.25 (2210 to 2210) was introduced in a PTF to MRNS Release Manufacturing and Delivery (ISMD) as well as being preloaded/shipped with current MRNS Release 2 orders.

2.2.8 The IBM 2210 as an IP Router

The IBM 2210 supports three dynamic routing protocols. All three routing protocols can run simultaneously on the IBM 2210.

The IP dynamic routing protocols supported by IBM 2210 are:

- Routing Information Protocol (RIP)
- Open Shortest Path First (OSPF)
- Exterior Gateway Protocol (EGP)

Additionally, the IBM 2210 implements IP multicasting routing protocols MOSPF and DVMRP.

The IBM 2210 supports ARP Subnet Routing (RFC 1027), also known as Proxy-ARP, and static routing.

This section describes the IP routing implementation on the IBM 2210.

The IBM 2210 implements the following IP functions:

- IP

 This is an unreliable and connectionless delivery mechanism which defines the IP datagram and specifies the delivery of these datagrams across the underlying network.

- ICMP

 Internet Control Message Protocol is used to report errors and provide information about unexpected circumstances. It includes support of Echo Request/Reply messages (known as PING), redirect messages (to direct a host to use another hop) and Source-Quench messages (used for congestion control).

- TCP

 Transmission Control Protocol is the connection-oriented protocol that allows the reliable stream delivery of data across a network from a TCP module on one machine to a TCP module on another machine.

- Telnet

 A simple remote terminal protocol that allows a user at one site to establish a TCP connection to a Telnet server at another site.

- UDP

 User Datagram Protocol provides a mechanism that allows application programs to send datagrams to other application programs.

- SNMP

 Simple Network Management Protocol is used to monitor IP routers and the network to which they attach.

- TFTP

 Trivial File Transfer Protocol is a simple file transfer protocol which runs on top of UDP.

- BOOTP

 The Bootstrap Protocol is used by diskless machines to learn their IP address and the location of the boot file and boot server.

2.2.8.1 General IP Parameters

When planning to use the IBM 2210 as an IP router, there are a number of IP parameters that you may configure regardless of the routing protocol used in your IBM 2210. These parameters are:

- Internal IP address

 You may assign an internal IP address to the IBM 2210. The internal address belongs to the router as a whole, and not to a particular interface. This address is always reachable as long as one interface on the router is active. This address is also used by the Data Link Switching (DLSw) feature.

- Router ID

 You may also assign a router ID to your IBM 2210. This is the default IP address used in various kinds of IP traffic originating from the router. For example, it is used as the IP source address in PING, TFTP or Traceroute packets.

- Routing table size

 Each IBM 2210 has a routing table which contains the dynamic routing information known by your router. Each entry in the routing table is 64 bytes, and, by default, the routing table size is 768 entries.

 You may change the number of entries in the IP routing table based on the requirements of your network.

- Router cache size

 The IBM 2210 uses a routing cache which contains the recently routed destinations. The router will reference the cache first before using the routing table. The minimum and default size for the router cache table is 64 entries. However, you may change the router cache size based on your requirements.

- IP broadcast format

 IBM 2210 allows you to specify the format that is used by your IBM 2210 when broadcasting packets out on a specific interface. In doing so, you must specify the style and the fill-pattern used.

 The style parameter can be either local-wire or network.

 When you specify local-wire for the style, the router will use the broadcast address of either 255.255.255.255 or 0.0.0.0. The former is used if you have specified the fill-pattern to be 1, and the latter is used with a fill-pattern of 0.

 When you specify network for the style, the router will send the broadcast messages that begin with the network and the subnetwork portion of the IP address of the interface. The host portion of the broadcast messages are either all 1s or 0s depending on the value specified for the fill-pattern parameter.

 Note: When receiving messages, the IBM 2210 recognizes all forms of the IP broadcast addresses regardless of the settings of these parameters.

- Reassembly size

 You can configure the size of the buffers that are used for the reassembly of the fragmented IP packets received by the router.

 By default, IBM 2210 uses a buffer of 12000 bytes.

 You can configure a route to a default gateway and the cost of reaching that default gateway. Normally, the default gateway is a router which has more routing information about the network.

- Default subnetwork gateway

 In a subnetted network, you can configure a separate default gateway and the cost of reaching it, for each subnet network.

 All of the packets detained for unknown subnets of a known subnetted network are forwarded to the subnetwork′s default gateway.

- IP access control

 The Access Control system allows the IBM 2210 to determine which packets are to be forwarded and which packets are to be discarded. For more information, refer to 2.2.8.10, “Access Controls” on page 111.

2.2.8.2 Interface Address Assignments

When you assign IP addresses to the router, you must note the following:

- You must assign at least one IP address to an interface. A hardware interface does not accept or send IP packets unless it has at least one IP address.
- It is possible to assign more than one IP address to an interface.
- You must specify an IP address together with its subnet mask.

Note

Serial lines do not need addresses. Such lines are called unnumbered and
can be configured without IP addresses, but you must still enable them for IP traffic using the following command:

```
IP Config>Add address 1 0.0.0.1
```

Using un-numbered serial lines has some restrictions which are documented in information APAR II08361.

2.2.8.3 RIP Implementation in IBM 2210

The following must be considered when configuring RIP for your IBM 2210:

- Only the network portion, as defined by a mask, is entered into the routing table.
- Masks are not sent in RIP broadcasts.
- Maximum number of hops is 15 and a hop count of 16 indicates infinity.
- Destination entries time out after three minutes.
- RIP updates are sent every 30 seconds.
- Variable length subnet masks are not supported.
- RIP is not supported across X.25 circuits.
- Split horizon is always used.
- Poison reversed may be enabled for individual interfaces.
- The 2210 does not accept host-routes in RIP updates.

RIP Interoperability with 6611 Network Processor: To use RIP between the 6611 Network Processor and the IBM 2210 you need to take the following into consideration:

1. The broadcast address type used by the IBM 2210.

 The 6611 only recognizes local-wire broadcasts. In our case, testing with V1R3 of MPNP, we found that both filling types are accepted. So broadcasting to 255.255.255.255 or 0.0.0.0 are both accepted by the 6611.

2. IBM 2210 does not accept host IP routes.

 The 2210 does not accept host-routes in a RIP response. The 6611 will advertise only the host address (not the network address) for the attached neighbors using the point-to-point protocol (PPP).

3. The RIP version configured for the 6611 Network Processor.

The 6611 Network Processor can be configured to use either RIP Version 1 or RIP Version 2. IBM 2210 only supports RIP Version 1. Therefore, when using RIP between the IBM 2210 and the 6611 Network Processor, the 6611 must be configured to use RIP Version 1.

2.2.8.4 OSPF Implementation

OSPF implementation sets the OSPF router ID to the address of the first OSPF interface appearing in the router's configuration. However, you may change the router ID using the configuration commands from the ASCII console or the General panel in the IP subdirectory of the Nways MRNS Configuration Program.

Note

When you change the router ID of your IBM 2210, the link state advertisements originated by the router before the router ID change may persist in the network for as long as 30 minutes. This may cause an increase in the size of link state database.

The OSPF implementation in the IBM 2210 provides support for TOS-based (Type Of Service) routing for TOS 0 only.

IBM 2210 provides support for simple password, allowing for the authentication of the link state advertisement received from the other routers. To provide authentication, you must do the following:

1. Specify authentication type 1 when you define the OSPF area.
2. Specify the authentication key to be used when you configure the OSPF parameters for each interface.

You can import routes learned from other protocols (EGP, RIP or static routes) into the OSPF domain when the OSPF router is configured as an AS boundary router. An OSPF router can also originate a default route into the area. For these purposes you need to enable AS boundary routing.

OSPF and Non-Broadcast Networks: If the IBM 2210 is connected to a non-broadcast multiaccess (NBMA) network and is eligible to become the designated router, you need to provide the router with the information to find its OSPF neighbor(s). You can achieve this by performing the following tasks:

- Define the interface to the NBMA network as non-broadcast.
- Specify the IP address of the OSPF neighbor(s) on the NBMA network.
- Configure your IBM 2210 to become the Designated Router.

In a star frame relay network with only 2210s, you can use the OSPF point-to-multipoint frame relay enhancement. Refer to Figure 78 on page 108 for an example of a star or partially meshed network. This type of network is also known as a spoke and hub network.

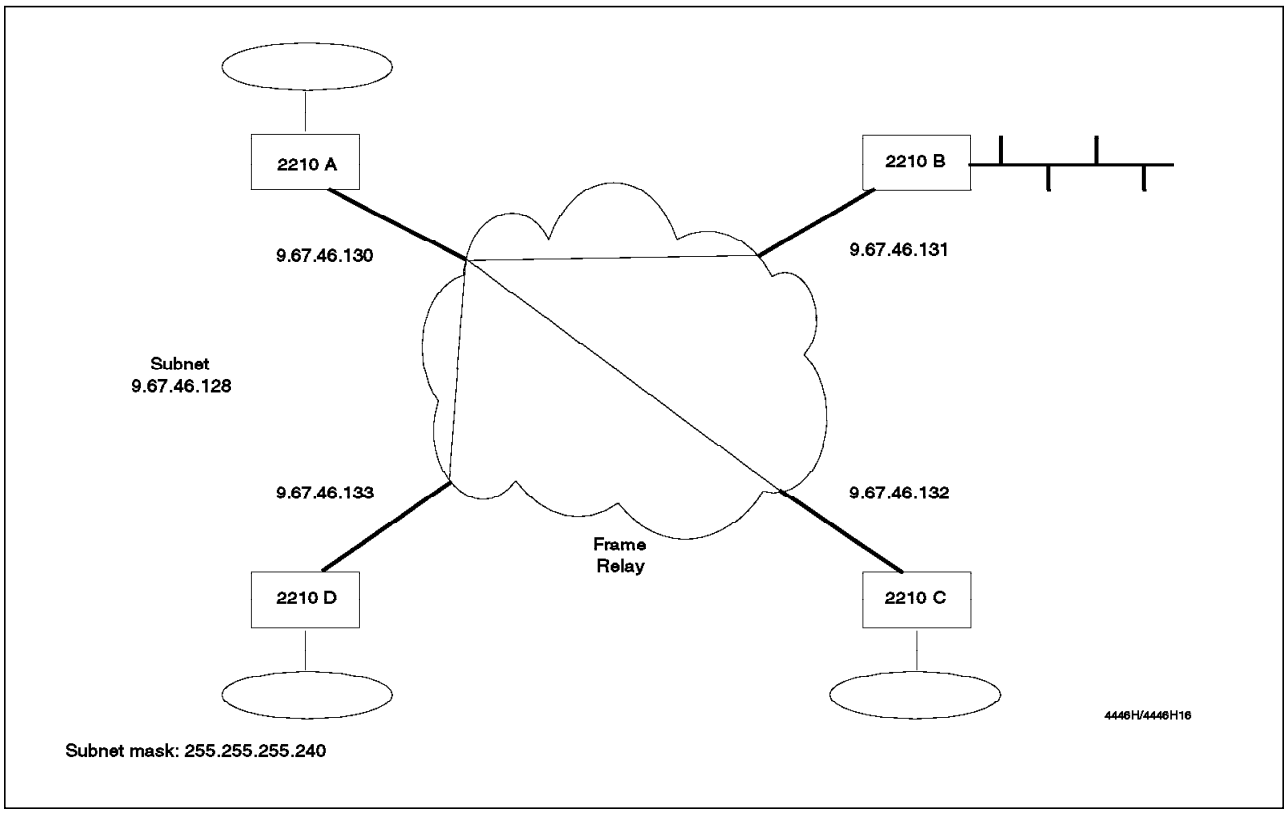


Figure 78. OSPF Point-to-Multipoint Frame Relay

Using the OSPF point-to-multipoint frame relay enhancement provided by IBM 2210, you may now assign a single IP subnet to an entire frame relay cloud and thus a single IP address to each frame relay interface of the router. In this case you only need to specify the OSPF neighbor at one side of each DLCI. In configuring such a network, you need to perform the following tasks:

1. Assign an IP address to the frame relay interface.
2. Enable OSPF on this interface.
3. Define the OSPF neighbor on one side of each DLCI (PVC).
4. To prevent one of the spokes from becoming the designated router, specify a router priority of 0 for the spokes and anything else but 0 for the hub router.

Note

In this type of OSPF configuration environment, it is not necessary to use the set non-broadcast command for each interface. By not using this command the router will determine that you intend to use the OSPF point-to-multipoint frame relay enhancement.

OSPF Interoperability with 6611 Network Processor: There are no specific OSPF considerations for connecting the IBM 2210 to the 6611 Network Processor when using OSPF.

Concerning frame relay, OSPF and 6611 interoperability, two scenarios were tested: scenario A and B.

- A: A fully meshed frame relay network with two 2210 routers and one 6611

- B: A partially meshed frame relay network in a star configuration where the 6611 is the hub and the 2210 routers are the spokes

Scenario A: Below, the steps concerning frame relay and OSPF are summarized, including the 6611 basic definitions:

- Assign an IP address to the 2210 frame relay interface.
- Enable OSPF and assign the interface to be an OSPF interface.
- Specify the interface as non-broadcast.
- Specify the 6611′s IP address as your OSPF neighbor on that interface and make it eligible to become the designated router.

On the 6611:

- Assign an IP address to the 6611 frame relay interface.
- Specify this interface as fully meshed.
- Enable OSPF and assign the interface to be an OSPF interface.

The interface type on the 2210 is multispecifying a nonbroadcast multiaccess (NBMA).

Scenario B: The differences are summarized in the steps below:

- Assign an IP address to the 2210 frame relay interface.
- Enable OSPF and assign the interface to be an OSPF interface.
- Specify the 6611′s IP address as your OSPF neighbor on that interface and make it eligible to become the designated router.

On the 6611:

- Assign an IP address to the 6611 frame relay interface.
- Specify the DLCIs with their destination IP address as point-to-point links.
- Enable OSPF and assign both interfaces, represented by the IP destination address, as the OSPF interfaces.

The interface type on the IBM 2210 is point-to-multipoint. Using this configuration, the spoke routers can still reach each other via the hub. The 6611 will take care of the routing between the spokes.

2.2.8.5 MOSPF

Multicasting is already used within OSPF. OSPF packets are sent to a standard multicast IP address of 224.0.0.5.

The 2210 extends this mechanism by implementing Multicast OSPF (MOSPF). When you enable the multicast forwarding capability, for each interface you can specify the following:

- Enable multicast forwarding on the interface.
- Enable the forwarding of multicast packets as unicast or multicast.
- Configure the IGMP polling interval.
- Configure the IGMP local database timeout.

The MOSPF function is used by the IBM 2210 for DLSw and IP Tunneling. Both implement client/server groups and peer groups for partner definitions.

DLSw uses a base multicast address of 225.0.1.0 for client and peers and an address of 225.0.65.0 for servers. The last octet of this address is used to identify the DLSw group number of the client/server group or peer group.

The IP bridge tunnel uses 224.168.0.0 as a base address for client/server groups as well as for peer groups. In this case the last two octets are used to identify a group.

Within this implementation it is also possible to manually change these addresses and to join or leave a multicast group specifying its IP address.

2.2.8.6 DVMRP

Distance-Vector Multicast Routing Protocol (DVMRP) allows you to define IP tunnels between MOSPF domains and a DVMRP domain/router. You can configure an IBM 2210 to use DVMRP and define interface(s) to use it.

2.2.8.7 EGP Implementation

EGP implementation includes the following:

- You can configure the set of routes you want to exchange with a particular neighbor by using the interchange flags and the interchange tables. In addition, you can select the cost you want to assign to a route.
- An EGP router may advertise itself as the default router via its IGPs (OSPF and RIP). This is called originating default. For information about specifying as a default router, refer to 2.2.8.3, “RIP Implementation in IBM 2210” on page 106 and 2.2.8.4, “OSPF Implementation” on page 107.

EGP Interoperability with 6611 Network Processor: There are no specific EGP considerations when connecting the IBM 2210 to the 6611 Network Processor.

2.2.8.8 Static Route Implementation

You can define a static route for:

- Default gateway

 Packets are routed to the default gateway when the destination cannot be found in the routing table.

- Default subnet gateways

 If you are using subnetted networks, you can define a separate default gateway for each subnetted network.

- Static network/subnet routes

 For each destination that is to have a fixed route, you can define a static route.

2.2.8.9 IP Filters

You can use IP filters to prevent forwarding of the packets for a network or subnet. This includes distribution of routing information about these networks.

2.2.8.10 Access Controls

The access control system allows you to be much more specific in filtering IP traffic. You can control access to particular classes of IP addresses and services by controlling source and destination IP addresses, IP protocol number and port numbers for the TCP and UDP protocols.

When you enable access control and add an entry to the list, all of the IP packets originated, forwarded, or received by the router are checked against the access control list. The following rules apply to this checking mechanism:

- For each packet received, the headers are compared to all of the specified fields in each entry in the list.
- If the entry matches the packet and the entry is inclusive, the packet is forwarded.
- If the entry matches the packet and the entry is exclusive, the packet is discarded.
- If there is no match with the entries in the access control list, the packet is discarded.
- Each entry has an IP address as well as source and destination IP address.
- Each IP address is logically ANDed with the mask and compared to the address in the entry.
- A mask of 255.255.255.255 matches only the resulting address itself.
- A mask of 0.0.0.0 and the resulting address of 0.0.0.0 is a wildcard and matches any IP address.
- Each entry may have an optional IP protocol number range. A range of 0 to 255 will match to all IP packets (within the address range).
- Each entry may have an optional port number range for UDP or TCP headers.

This implication of the above rules is that if you want to make one exclusion, you need to add inclusion(s) for all of the other IP traffic you want to be forwarded by the router.

2.2.8.11 BOOTP Implementation

The IBM 2210 implements the Boot Process (BOOTP) Client function and the Boot Process (BOOTP) Relay Agent also known as BOOTP Forwarder. The 2210 may use the BOOTP client function to obtain its boot file (refer to 2.2.7.1, “Boot Files and Boot Processes” on page 94). It may also be configured to forward BOOTP requests to a BOOTP server.

The 2210 cannot act as a BOOTP server. You need a host running the BOOTP daemon. A BOOTP server contains a file that lists all of the BOOTP clients for which this server is responsible for, their associated IP addresses, and the location and name of their boot files.

The following is a summary of the BOOTP process:

1. The BOOTP client copies its MAC address into a BOOTP packet (based on UDP) and broadcasts it onto the LAN.

2. If the BOOTP client and server are not on the same network, a local BOOTP relay agent will receive the request from the client, and route it to its

defined BOOTP server(s) or to the next BOOTP relay agent and route to the BOOTP server.

3. The BOOTP server receives the request and tries to match the MAC address with one in its list. If it finds a match, it will send a BOOTP reply with the client's IP address, subnet mask, and BOOTP server name. If the BOOTP client and server are not on the same network, the BOOTP reply may go through relay agent(s) to reach the client. In this case, the relay agent will receive a BOOTP reply, add an entry to its ARP table and forward the reply to the client.

4. The client uses the information that is contained in the reply to initiate a TFTP request to the TFTP server to download the boot image.

You need to assign two parameters when you define the router as a BOOTP forwarder (relay agent):

- The maximum number of hops you want the BOOTP request to go through. This is not the number of IP subnetworks, but the number of BOOTP relay agents needed to get the server from the client (and vice versa).
- The number of seconds you want the client to retry before the BOOTP request is forwarded. BOOTP uses a technique of timeout and retransmission. When a client sends a BOOTP request, it starts a timer. If it does not receive a response before the timer expires, it retransmits the request. This process will be repeated the number of times that you have specified.

2.2.8.12 Telnet Implementation

To allow you to access the ASCII console interface remotely, the IBM 2210 implements the Telnet function. It allows you to have five Telnet sessions: two servers (inbound to the router), and three clients (outbound from the router).

The Telnet session to the IBM 2210 does not provide you with any indication of which router you are logged into. You may determine the router by displaying the configuration information of the router. Alternatively, you may use Ctrl+Break to access the Telnet command mode. You can then issue the `status` command to display the IP address of the station that you are connected to as well as the current terminal mode.

2.2.8.13 SNMP Implementation

Simple network management protocol (SNMP) runs on top of the user datagram protocol (UDP) and is used for monitoring and managing IP hosts in an IP network. SNMP enables network hosts, running vendor-supplied software, to read and modify some of the router's operational parameters. In this way, network management is established for the IP community. The software that processes the SNMP requests from the network management hosts runs on the IBM 2210 and is called an SNMP agent.

The following are the various aspects of the SNMP that you need to consider when configuring the SNMP for your IBM 2210.

Authentication: In SNMP you can define a community. The SNMP community is simply a group of nodes that share network management information. The community is established at configuration time.

The community allows you to define the IP address of the SNMP management station that is allowed to access the information in the SNMP agent's Management Information Base (MIB). It allows you to define a community name in accessing the MIB. The community name is used as an authentication scheme that prevents unauthorized users from learning information about an SNMP agent or modifying its characteristics. By defining an authentication scheme, you can provide security in your network management system.

Note: If no IP address is defined for the SNMP manager in your community table, any IP station that provides the correct community name will be able to access the MIB in the SNMP agent.

MIB Support: The operational parameters or variables are defined by a MIB. The standard MIBs supported by IBM 2210 are described in Appendix D of *The Nways MRNS Protocol and Monitoring Reference*.

For each community name, you can specify which MIB or which part of a MIB can be accessed by the members of that community. To do so, you must first add one or more MIB Object IDs (the identification of a MIB item) to a view, creating a sub-tree. Then you assign a view to a community.

Traps: SNMP agents can create trap messages. These are unsolicited messages that are sent from the router to an SNMP manager in response to a router or network event or condition, such as a router reload or network down. The IBM 2210 provides two types of traps which can be enabled or disabled separately for a specific community name:

- General traps

 These traps are defined by the RFCs and allow the router to send the traps asynchronously to the SNMP manager in case of a specific event. There are six general traps defined:

 - Link-up
 - Link-down
 - Cold start
 - Warm start
 - EGP neighbor loss
 - Authentication failure

- Enterprise-specific traps

 These traps are specific traps which can be generated by event logging system (ELS) messages. You can use the ELS trap command to enable sending of messages or groups of messages via an SNMP trap. To enable this to be forwarded by the SNMP agent of your router, you need to enable the trap type enterprise. However, the SNMP manager must support these enterprise traps because they are specific to the IBM 2210.

2.2.8.14 TFTP Implementation

The IBM 2210 implements the TFTP client function and the TFTP server function. The client function allows you to send or receive configurations or boot images to and from a TFTP server. The server function is implemented to provide other routers with a boot image or a configuration file. This implementation allows multiple, simultaneous file transfers between the router's nonvolatile configuration memory (NVCNFG), the Integrated Boot Device (IBD), and remote

hosts. Refer to 2.2.7.1, “Boot Files and Boot Processes” on page 94 for more information about the boot mechanism.

The TFTP implementation does not allow you to use PUT or COPY to transfer files to another router.

When a router acts as a TFTP server, transfers are transparent to the user. Use the ELS message log to view the transfers in progress. To view all TFTP messages, go to the ELS prompt of the GWCON and issue the following commands:

```
+ event

ELS>display subsystem tftp all
```

You can view the messages by using either of the following commands: displays the messages on the CONFIG console:

divert 2 0 Displays the messages on the CONFIG console

talk 2 Displays the messages on the MONITOR console

2.2.8.15 ARP Subnet Routing

The IBM 2210 implements Proxy-ARP router function. When the router is configured for ARP subnet routing, it will reply by proxy to the ARP requests for destination which are reachable via the 2210′s interfaces.

2.2.9 Data Link Switching

This section provides a brief overview of data link switching (DLSw) and discusses configuration of data link switching on the IBM 2210.

2.2.9.1 Data Link Switching Overview

DLSw is designed to facilitate integration of SNA traffic into a multiprotocol network. DLSw functions include:

- Transporting of SNA in a multiprotocol routed backbone
- Dynamic rerouting in the wide area network
- Reliable delivery of SNA traffic
- Termination of LLC acknowledgements on the LAN segments
- Broadcast traffic control through the WAN
- LAN and WAN control for congestion and data flow

DLSw uses IP encapsulation of SNA as its transport vehicle across the internetwork. To supply the reliability SNA requires in the internetwork, DLSw uses Transmission Control Protocol (TCP) flows between edge-node routers (those routers joining the LAN segments to the IP portion of the network).

DLSw routers establish TCP connections to other DLSw routers using ports 2065 and 2067. Port 2065 is a read port on which all DLSw information is received, and port 2067 is a write port from which all DLSw information is sent.

DLSw also uses a technique known as DLC termination, or spoofing, to minimize T1 timer expirations and to keep acknowledgements isolated to the local LAN segment.

Spoofing is the process that acknowledges receipt of the frame on the local LAN segment by masquerading as the destination end station. Spoofing keeps the receiver ready and/or supervisory poll frames from leaving their subnet media. Therefore, it ensures local media response speeds to acknowledge layer 2 timers (T1 timers for example) and lessens the bandwidth overhead requirements in the WAN.

2.2.9.2 DLSw on the IBM 2210

The DLSw function of the IBM 2210 supports the interconnection of SNA devices attached to either a LAN (token-ring or Ethernet) or an SDLC multipoint non-switched line.

As a prerequisite for DLSw, if the IBM 2210 supports LAN-attached SNA devices, it must be configured to support source-route bridging on the token-ring interface, or transparent bridging on the Ethernet interface.

A DLSw virtual segment number also needs to be configured for IBM 2210s implementing DLSw. This virtual segment must be the same for all IBM 2210s participating in the DLSw function. This is to ensure that the end stations both see the TCP/IP network as one token-ring.

SNA devices attached to an IBM 2210 via SDLC multipoint non-switched lines are each assigned a token-ring locally administered address (LAA), service access point (SAP) and SNA XID (Exchange ID). These will be used by the IBM 2210 to represent such devices to other SNA devices that are using the DLSw function as if they are attached to a token-ring LAN. SDLC-attached devices can have SNA connections with token-ring and/or Ethernet-attached devices connected to the same IBM 2210.

SNA devices attached to an IBM 2210 establish connections with SNA devices attached to other IBM 2210s as if they are on the virtual segment.

SNA devices attached to an IBM 2210 via LAN segments establish connections with SNA devices attached to the same IBM 2210 via SDLC as if they were on the virtual segment.

Data Link Switching Supported Topology: There are two types of data link switching:

- Local data link switching
- Remote data link switching

In local DLSw, the data link switching function is performed within a single IBM 2210. In remote DLSw, stations attached to two or more IBM 2210s communicate across an IP network using DLSw.

Local Data Link Switching: Local DLSw allows communication between a token-ring or Ethernet-attached SNA device and an SDLC secondary PU2.0 or PU2.1 station that is link attached to the IBM 2210.

With Version 1 Release 2 of the IBM 2210 Nways MRNS software, both PU2.0 and PU2.1 link stations can coexist over SDLC lines at the same time.

The LAN-attached device is locally attached to the same IBM 2210 or attached to a remote LAN which is bridged to your IBM 2210.

Each SDLC-attached PU2.0 or PU2.1 device is assigned a MAC and SAP address and will appear to the other SNA devices as if it is attached to a token-ring LAN on your IBM 2210. Local DLSw converts SDLC frames to LLC2 frames. The encapsulated SDLC frames are passed to the DLSw function which will in turn use the source-route or transparent bridging function to deliver them to the LAN-attached device.

Remote Data Link Switching: SNA stations attached to an IBM 2210 via a token-ring, Ethernet or SDLC connection can establish sessions with other SNA stations which are attached to a remote IBM 2210 or 6611 Network Processor via a token-ring or an Ethernet connection. The connection between the two IBM 2210s or between the IBM 2210 and the 6611 Network Processor is over an IP network which can include OEM routers which support compatible IP functions such as RIP or OSPF. Note that only the two routers connected to the end stations must be enabled for DLSw. The DLSw function is not required in the routers which might exist between the two edge-node routers.

The DLSw in the IBM 2210 encapsulates the SNA frames in a TCP/IP datagram and delivers the encapsulated frames to its partner over the IP network.

Remote DLSw supports:

- SDLC to LAN over WAN

 SDLC frames are converted into LLC2 frames. This allows a link-attached SDLC secondary device to communicate with a LAN (token-ring and Ethernet) attached device.

- LAN to LAN over WAN

 Remote DLSw allows communication between SNA devices attached to token-ring or Ethernet networks. Remote DLSw can convert frames between the token-ring and Ethernet allowing token-ring and Ethernet-attached devices to communicate with each other using DLSw.

DLSw Using MOSPF: The IBM 2210 supports the use of the DLSw Group Membership function to allow it to dynamically discover its DLSw partners, instead of having to manually configure the partner addresses. This feature utilizes the Multicast OSPF (MOSPF) function, which is described in 2.2.8.5, “MOSPF” on page 109.

The DLSw Group Membership defines two types of groups:

- Client-to-server
- Peer-to-peer

Client-to-server groups have members that are designated as either a client or a server. Server routers only form DLSw connections with client routers. This group type is used for subarea SNA connections. Peer-to-peer groups have members that are all designated peers. All members of a peer-to-peer group will form DLSw connections with all other members of the group. This group type could be used for APPC connections.

DLSw group membership will only work between routers that support it, so a combination of group membership and preconfigured DLSw partner definitions may be required in your network.

2.2.10 Features and Facilities

This section describes the different features provided by the IBM 2210, the Bandwidth Reservation (BRS), the MAC Filtering (MCF), and the WAN Restoral (WRS) also called Dial Backup. It also describes some facilities provided by the IBM 2210 such as the dial-on-demand, NetBIOS name caching, and NetBIOS filtering.

2.2.10.1 Bandwidth Reservation (BRS)

In this section, we explain the Bandwidth Reservation feature, we show the Bandwidth Reservation configuration commands, and a scenario of Bandwidth Reservation is provided.

Introduction to Bandwidth Reservation (BRS): The Bandwidth Reservation feature allows you to reserve part of the bandwidth on the link for a specific traffic type.

Note:

- For Version 1 Release 1 of the Nways MRNS software for the IBM 2210, Bandwidth Reservation (BRS) is supported only over PPP serial links and applies to outbound traffic only.
- For Version 1 Release 2 of the Nways MRNS software for the IBM 2210, Bandwidth Reservation (BRS) supports the point-to-point protocol, frame relay, and dial circuits (ISDN and V.25 bis). Again this applies to outbound traffic only.

Figure 79 shows specific data streams assigned to a part of the WAN bandwidth.

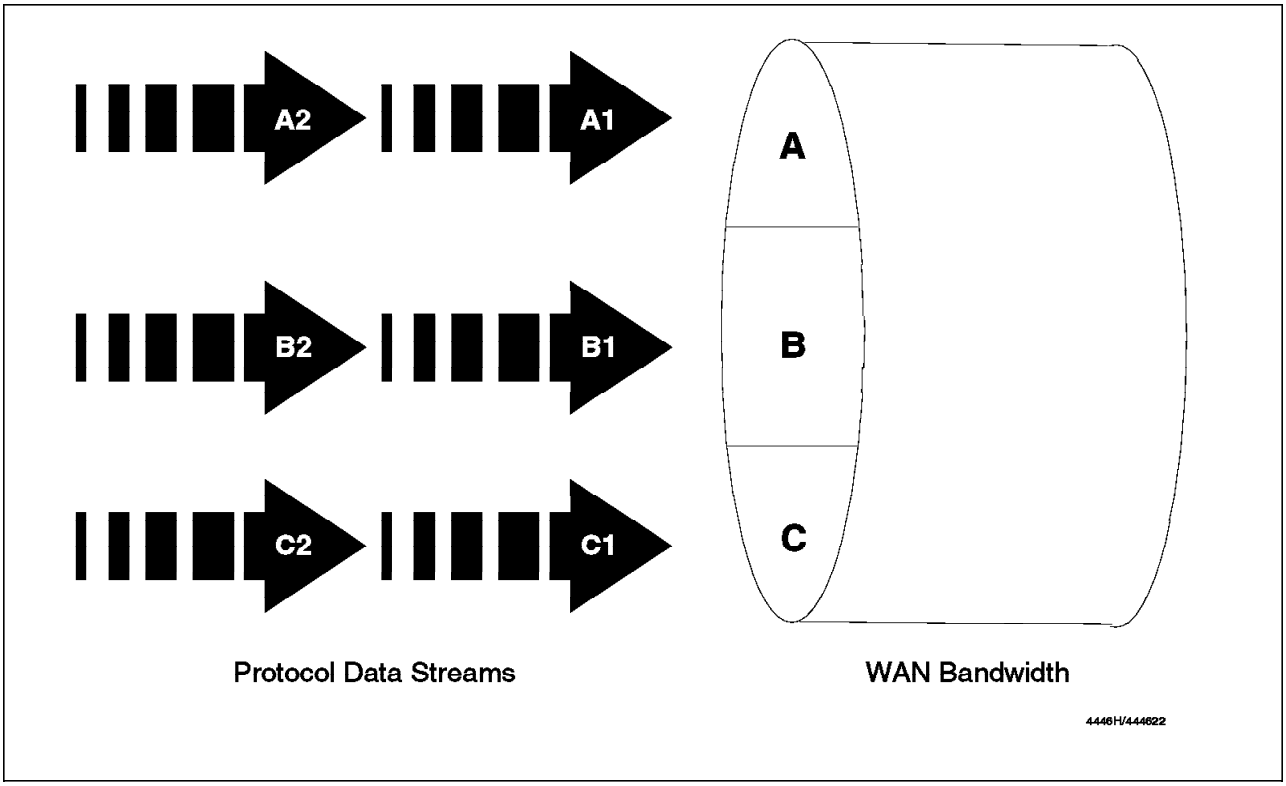


Figure 79. Bandwidth Reservation

First of all, you assign a name to a percentage of the bandwidth. This is called a class name.

Note: All the names of the classes are case sensitive.

By default, there are two classes of names that you can neither delete nor change. You are only allowed to change their percentage of the bandwidth. These two classes, by default, are:

- LOCAL with 10% of the bandwidth by default
- DEFAULT with 40% of the bandwidth by default

The total of all the percentages of all the classes defined must *not* exceed 100%.

The reserved percentages are the guaranteed minimum slice of the bandwidth for the network connection. If the network is operating at full capacity, the messages from a specific traffic class can only be transmitted as long as they don't use more bandwidth than allocated for that class. If the rate of the messages exceeds the reserved bandwidth, the messages are held until other bandwidth transmissions have been satisfied.

In the case of light traffic on the network, packet streams can use bandwidth exceeding their allowed minimum (up to a maximum of 100% of the bandwidth) if there is no other traffic.

When you assign a class to a type of traffic, you must also assign the priority class of this traffic within its class. There are four priority classes:

- Low
- Normal
- High
- Urgent

For example, a traffic assigned with class DEFAULT and priority *urgent*, will be delivered faster than a traffic assigned with class DEFAULT and priority *normal*.

The priority setting within the bandwidth class has no effect on other bandwidth classes. That is, none of the bandwidth classes have priority over the others.

Note: If no priority is assigned within a class, the default priority is *normal*.

After defining the class names, you may assign these classes to the following traffic types:

- The DEFAULT traffic class

 The DEFAULT traffic class is used by all the traffic that is not assigned to a specific class. By default, the DEFAULT traffic class uses the class DEFAULT, with the default class priority *normal*.

- The protocols (IP, ARP, IPX, ASRT, APL or AP2)

 For protocols, you can assign a specific class and priority for each of the following protocols:

 – IP

 – ARP (with ASCII console only)

- IPX
- ASRT (Means bridged traffic)
- APL (AppleTalk phase 1)
- AP2 (AppleTalk phase 2)

Note: The ARP protocol is not currently available on the Nways MRNS Configuration program. You must customize it via the Nways MRNS program on the ASCII console.

- The *filter* (RLOGIN_IP, TELNET-IP, NetBIOS, SNA Bridged, SNPM-IP, DLSw-IP, MULTICAST-IP, TUNNELING-IP and SDLC-IP)

 For the filters, you can assign a specific class and priority for each of the following filters:

 - RLOGIN_IP
 - TELNET-IP
 - NetBIOS (bridged NetBIOS traffic)
 - SNA (bridged SNA traffic)
 - SNMP-IP
 - DLSw-IP (SNA traffic via DLSw)
 - MULTICAST-IP
 - TUNNELING-IP (with ASCII console only)
 - SDLC-IP (with ASCII console only)

 The TUNNELING-IP filter and the SDLC-IP filter are not currently available on the Nways MRNS Configuration program. You must customize them via the Nways MRNS program on the ASCII console.

- Five TAGs (from MAC filtering on bridged traffic only)

 You can assign a specific class and priority for the following tags defined by the MAC Filtering (MCF) feature:

 - TAG1
 - TAG2
 - TAG3
 - TAG4
 - TAG5

 Note: The TAG number is assigned to a bridged traffic with the MAC filtering features.

2.2.10.2 WAN Restoral (WRS)

This section provides a description of the WAN Restoral feature and its configuration commands. A scenario of how to configure WAN Restoral on the IBM 2210 is also provided.

Introduction to WAN Restoral (WRS): The WAN Restoral (WRS) feature, which is also called the Dial Backup feature, allows you to back up a primary leased PPP serial link with a switched V.25 bis PPP serial link.

Note

Backing up of frame relay or X.25 serial link is not supported. WAN Restoral only supports backing up of PPP leased serial link.

The WAN Restoral feature is supported over every routed protocol (IP, IPX, AppleTalk and DLSw) and for every bridging method, including tunnel bridge.

The backup switched line supported by this feature is over V.25 bis modem. In a future release, the WAN Restoral with the backup serial line over ISDN serial line will be provided for IBM 2210 models 127 and 128.

When the IBM 2210 detects the loss of connectivity on the primary PPP serial link, it automatically dials the configured phone number to establish the dial connection via the V.25 bis modem.

There is only one remote phone number configured in the IBM 2210. This must be the phone number of the same remote IBM 2210 which is reached via the primary serial link.

When the switchover from the primary link to the backup link occurs due to the failure of the primary link, the whole set of protocols configured on the primary leased PPP serial link will be automatically switched over to the switched V.25 bis serial link. All of the protocols (IP, IPX, AppleTalk, DLSw) and all of the bridging methods will survive the switchover to the switched V.25 bis serial link.

When the IBM 2210 detects that the primary PPP serial link has come back up, it automatically drops the V.25 bis dial connection and restores all the protocols to use the primary leased PPP serial connection.

Figure 80 shows the typical configuration of a network using WAN Restoral.

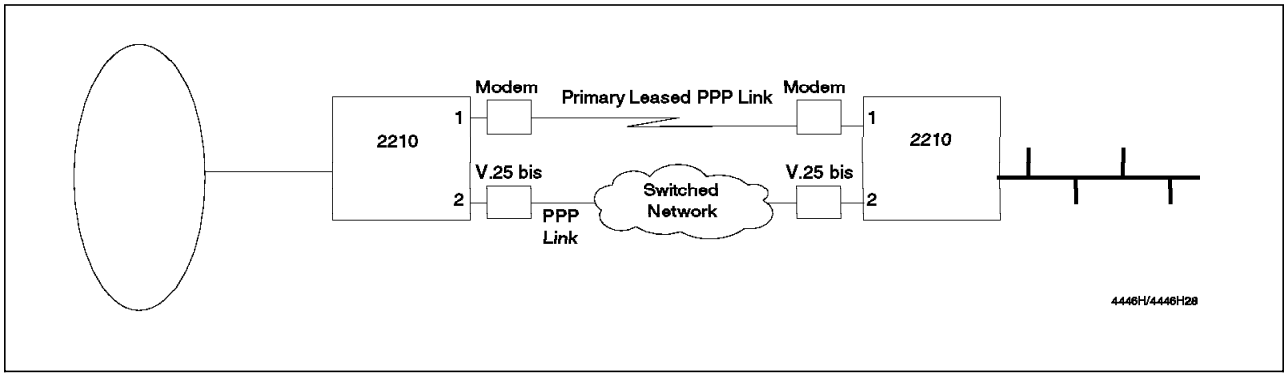


Figure 80. Typical Implementation of WAN Restoral

To be able to use the WAN Restoral, both 2210s at each end of the primary serial link must be customized for WAN Restoral.

To configure a 2210 to use WAN Restoral, you must customize one of its serial interfaces with the PPP link, and the other serial interface as a dial interface using the V.25 bis modem with the PPP encapsulation method.

Since this feature is not supported by the 6611 Network Processor, the only possible way to use this feature in a network that includes the 6611 Network

Processor is shown in Figure 81 on page 121. In this configuration, the IBM 2210 could detect the primary link failure and dial the 6611 Network Processor over the backup link.

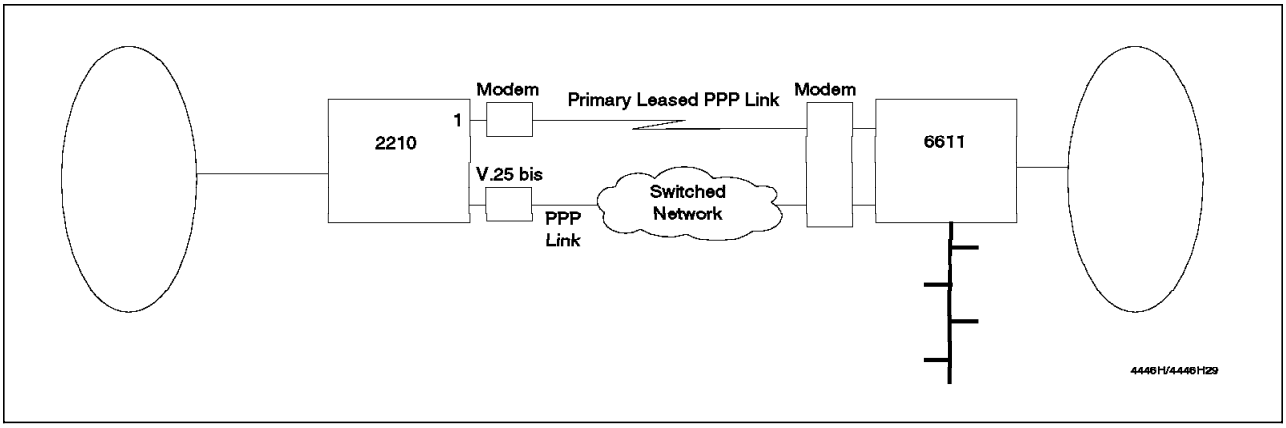


Figure 81. Possible Interoperability of WAN Restoral with 6611 Network Processor

2.2.10.3 Dial-on-Demand

This section provides a description of the dial-on-demand facility. It shows the dial-on-demand configuration commands and provides an example scenario of the dial-on-demand configuration.

Introduction to Dial-on-Demand: The dial-on-demand facility is designed for remote sites that do not need to be connected to the central site all of the time but only when there is some data to be sent.

When the IBM 2210 detects that a packet needs to be sent over the switched network to a remote IBM 2210, it automatically dials the customized phone number to establish the dial connection via the V.25 bis modem.

You could customize several phone numbers in the IBM 2210, and map each remote phone number to a specific protocol address (IP or IPX address). However, note that only one connection to a remote site is allowed at any single point in time. This means that if there is already a connection to a remote site, you cannot send any packets to another remote site. In this case, you must wait until the first connection is terminated before trying to reach the second remote site.

To use the dial-on-demand facility, you must configure all the parameters of the desired protocol (IP or IPX) on the corresponding virtual dial-circuits and not on the physical V.25 bis interface.

When the IBM 2210 detects that no more packets are required to be sent over the switched interface for a certain lapse of time (idle time), the switched line is automatically dropped and the V.25 bis modem becomes available.

Note that when you customize a serial interface as a dial interface using the V.25 bis modem with the PPP encapsulation method, the other physical serial interface is able to be used for anything else at the same time. Also, both 2210s at each end of the primary serial link must be customized for dial-on-demand.

Note

It is recommended that you allow only one site to issue outbound calls, and the other site should allow inbound calls only. This will prevent dial collision in case both sides want to call each other at the same time. However, this is not a requirement and you can enable both sides for both inbound and outbound calls. In this case, you must be aware that if the IBM 2210s want to call each other at the same time, the V.25 bis modems will loop with DIALING, then BUSY, then DIALING, then BUSY, etc. This will be repeated until one side decides to no longer send data to the other side. Then the switched link will be activated from the other side.

For IP routing over dial-on-demand, it is recommended that you customize static routes. This prevents the IBM 2210 from establishing the connection for each routing table update which is sent by the dynamic routing protocols.

If there is DLSw customization over a dial-on-demand circuit, be sure to not enable the Keepalive parameter. By enabling this parameter to verify that the remote DLSw partner is alive, the dial-up connection would remain active permanently.

IPX does not provide static routing. Therefore, you are advised to specify large RIP and SAP update intervals to ensure that the dial-on-demand circuits are not frequently established as a result of the frequent RIP and SAP messages in an IPX environment.

Note

Dial-on-demand cannot be used to provide additional bandwidth over a switched serial interface in case of overutilization of the bandwidth of a primary leased serial interface.

The dial-on-demand facility is only supported over:

- TCP/IP (including DLSw and Tunnel Bridge)
- IPX protocol

Note

Dial-on-demand is not supported for any bridging methods except for the tunnel bridge method which is actually using the IP protocol over the serial links.

Dial-on-demand is only supported over a switched V.25 bis PPP serial link.

Figure 82 on page 123 shows you a typical drawing of a dial-on-demand network.

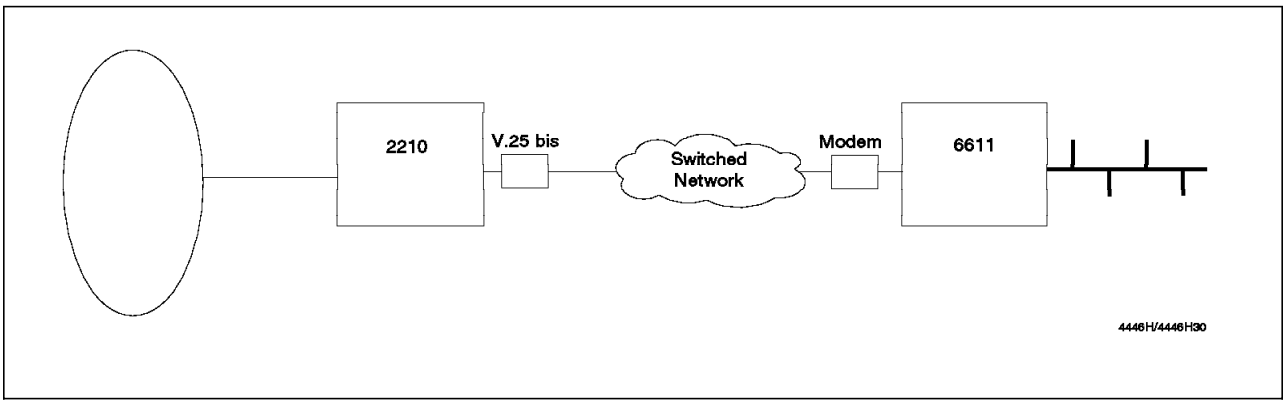


Figure 82. Typical Implementation of Dial-on-Demand Processor

This facility is not supported by the 6611 Network Processor; therefore, the only possible way of using this feature in a network which includes the 6611 Network Processor is shown in Figure 83. In this configuration, the IBM 2210 could dial the 6611 Network Processor when it has data to send to the 6611. But if the switched link is not up and the 6611 Network Processor has to send data to the IBM 2210, it must wait until the 2210 establishes the call. This will happen when the IBM 2210 has data to send to the 6611 Network Processor.

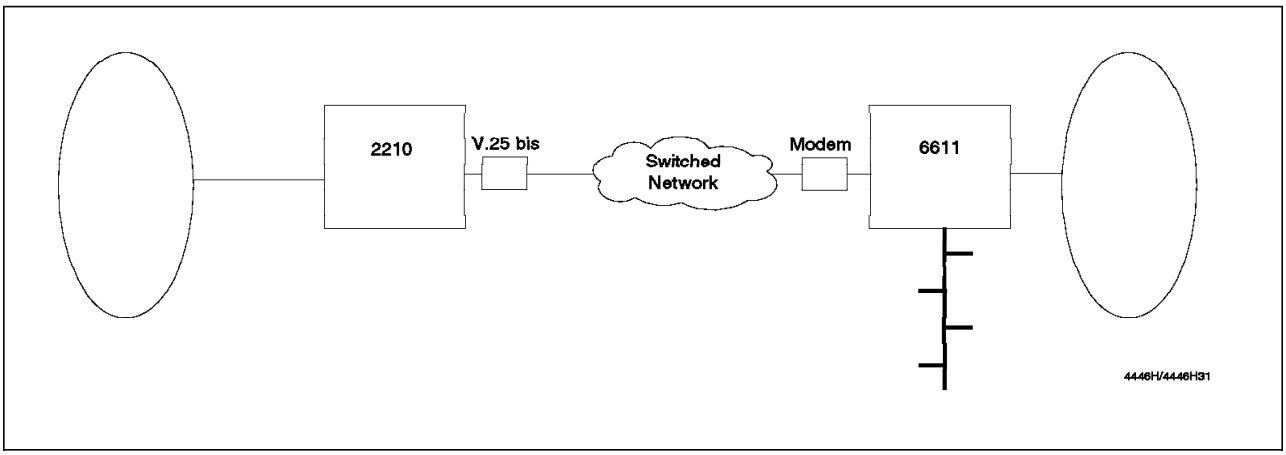


Figure 83. Possible Interoperability of Dial-on-Demand with 6611 Network Processor

2.3 IBM 6611 Router

This section provides a summary of the hardware and functions of the IBM 6611 Network Processor when used with the IBM Multiprotocol Network Program.

Further information on the IBM 6611 Network Processor hardware can be found in the *IBM 6611 Network Processor - Installation and Service Guide*.

Further information on the functions provided by the IBM 6611 Network Processor when used with the Multiprotocol Network Processor can be found in the *IBM 6611 Network Processor - Introduction and Planning Guide*.

The IBM 6611 uses its bridging, routing and data link switching functions to receive and transmit multiple protocols from one LAN to another. The Multiprotocol Network Program provides the necessary configuration functions to support each protocol. The 6611 is not a gateway and therefore requires the end

stations that want to communicate with each other to use the same protocol. The data link switching function encapsulates SNA and NetBIOS frames into an IP datagram for transport over a WAN. With all other protocols it uses the packet or frame format prescribed by that protocol to route or bridge that protocol. Each of the adapters has its own high-performance processor and is called a peer-capable adapter. Except in the case of data-link switching, the adapter processors eliminate the need to pass packets to the system processor enabling faster system performance and packet transfer.

The Multiprotocol Network Program collects and stores status information about the IBM 6611 connections. Performance and other data are stored in its MIB variables. Traps are sent to the SNMP manager for events that occur in the network and router itself. The SNMP manager can then retrieve MIB information to help with problem determination.

The 6611 supports local or remote access and control via the System Manager component of the Multiprotocol Network Program. This program allows you to set passwords, run software and hardware diagnostics, view statistics and error logs and shut down the 6611. Access can be via a local or remote interface.

2.3.1 Hardware Overview

There are three main user components that make up the 6611:

- The IBM 6611′s family
- The Multiprotocol Network Program (MPNP)
- The System Manager

There were many modifications to the IBM 6611′s family, as described below:

- New 6611 Model 120 configurations
- New 6611 Model 125
- New 6611 Models 145 and 175 replacing Models 140 and 170 respectively
- New adapters

2.3.1.1 Model 120 Enhancements

The following is a complete list of the Model 120 fixed configuration which will be available. The new configurations are:

- Four SDLC ports / two multi-interface serial ports
- One token-ring port and one Ethernet port
- Two token-ring ports
- Two Ethernet ports

The existing configurations are:

- One token-ring port and four SDLC ports
- One Ethernet port and four SDLC ports
- One token-ring port and one X.25 port
- One Ethernet port and one X.25 port
- One token-ring port and two multi-interface serial ports
- One Ethernet port and two multi-interface serial ports

The benefits with these changes are:

- **Expanded Configuration Options**

 These key new configurations will allow the 6611 to be used as a local bridge, both between like media as well as between disparate media. When used in conjunction with MPNP V1R3′s new Translational Bridging function, the 6611 Model 120 can now provide translational bridging between token-ring and Ethernet LANs.

- **Current Configurations Enhanced**

 The existing Model 120 configurations have been replaced by new configurations which utilize the new 6611 adapters, providing the improved performance and increased connectivity previously described. Even though the new 6611 adapters increase the number of ports per adapter, the Model 120s will still be limited to the same number of ports as today. In other words, if a combination adapter is used to achieve a configuration that is currently available on the Model 120, then the second adapter slot will not be used.

 For example, the one token-ring port and two multi-interface serial ports Model 120 configuration will now be handled by one adapter. The performance of the new Model 120 will be equivalent to the old Model 120 with the two adapters.

 The Model 120 configurations involving a four-port SDLC adapter or an X.25 adapter will use both slots of the Model 120. The other configurations will use the new adapters.

IBM 6611 Model 120 is positioned for the small or remote office with two LAN attachments.

2.3.1.2 IBM 6611 Model 125

This open, two-slot model complements the Model 120′s fixed configuration offerings. This versatile new model provides the following benefits:

- **Flexible configurations**

 The Model 125 can support any of the wide range of new 6611 adapters up to a maximum of eight ports. In many instances, the Model 125, coupled with the new multiport and combination adapters, can support a configuration which previously required a four-slot Model 140, representing a significant savings.

- **Future flexibility**

 Unlike the Model 120, which is available only in fixed configurations that cannot be changed after installation, the Model 125 gives customers the ability, in the future, to change adapters as their network configuration needs change.

 Adapters ordered for a Model 175/145 can be installed and used successfully in a Model 125. This allows flexibility in using adapters as the network needs change.

Note

Please be aware that adapters ordered for a Model 125 cannot be used in a Model 175/145. If a Model 125 adapter is installed in a Model 175/145, the adapter is marked as invalid at IPL time. When a configuration is attempted to be loaded into the 6611, the configuration will be invalid since the adapter is invalid.

- **Full function**

 While the Model 125 is a relatively small box in terms of the number of adapters supported, it is supported by the same software as the larger 6611 models with no restriction on the available functions.

IBM 6611 Model 125 is also targeted at the small or remote office, but it can handle up to three LANs and a couple of WANs.

2.3.1.3 IBM 6611 Models 145 and 175

As replacement models for the current Models 140 and 170, the Models 145 and 175 were designed to offer improvements in packaging and usability while maintaining the same external interfaces. In this way, customers can capitalize on the improvements provided while investing a minimal amount of time familiarizing themselves with the new models. The IBM 6611 Models 145 and 175 use the same physical environment.

These new four- and seven-slot models support any mix of the new adapters and offer the following benefits:

- **Rack mount options**

 There are two rack mount features available for the Models 145 and 175. One is a set of brackets that attaches to the sides of the box and permits installation on any industry-standard 19-inch, two- or four-rail open or closed rack (including the IBM 9309). This enables the optimal use of the space in wiring closets and machine rooms.

 If faced with installing a 6611 in an area which is densely populated with equipment or is in a hard-to-reach location customers may want to consider the sliding shelf feature. This exceptionally sturdy steel cantilevered shelf mounts on any industry-standard 19-inch rack and is equipped with a recessed handle which enables the shelf to be easily pulled forward, extending it to a depth of 27 inches. When the 6611 is placed on the shelf, the user has full range of access to all sides of the machine, significantly simplifying installation and removal of adapters or other maintenance activities. The 6611 can be screwed into the shelf, and the rubber feet sit in holes on the shelf to prevent the shelf from slipping.

- **Customer setup**

 The new models of the 6611 are designed to support customer setup, further streamlining the installation process. The new adapter features also support customer setup on the new models, making any future configuration changes easier to accommodate and schedule.

- **Space savings**

 The seven-slot Model 175, like the four-slot Model 145, is designed for either horizontal installation on a rack or used stand-alone on a table or desktop. This represents a considerable space savings compared to its predecessor, the Model 170, which could be installed only in a vertical position. The Model

175 is also considerably lighter, weighing only 42 pounds fully populated, compared to the Model 170′s maximum weight of 88 pounds.

- **Usability improvements**

 To enable easier access for attachment of an ASCII display or SCSI tape drive for diagnostics or service, the S1 service port and SCSI port have been moved to the front of the box. This makes cabling between the devices easier, as well as reduces the risk of disturbing an installed adapter cable or power cord.

 A cable management bracket is provided as a standard feature for both Models 145 and 175. This bracket mounts on the rear of the box to provide strain relief for adapter cables, as well as improve cable management by allowing each cable to be dressed through an individual opening.

- **External interfaces preserved**

 Although the packaging of the new models has changed, the interfaces that customers use have remained the same as the predecessor models. Models 145 and 175 use the same three-character display on the operator panel for information and error codes, support the function-rich System Manager for diagnostic and management tasks, and utilize the easy-to-use 6611 Configuration Program for initial and subsequent configurations. Use of these common configuration and management tools across the product line simplifies network operation and management, and protects customer investment in training and support resources.

- **Scalability**

 In the event that a change in a customer′s network configuration causes the requirements to exceed the capacity of the installed Model 145, a Model Upgrade is available to convert the Model 145 to a Model 175, enabling the use of three additional adapter slots.

 As network needs change, the adapters from Models 175/145 can be moved to another Model 175, 145 or 125. This allows flexibility in using adapters as network needs change.

 Note

 Please be aware that adapters ordered for a Model 125 *cannot* be used in a Model 175/145. If a Model 125 adapter is installed in a Model 175/145, the adapter is marked as invalid at IPL time. When an attempted configuration is to be loaded into the 6611, the configuration will be invalid since the adapter is invalid. Also, the old adapters for the Models 140 and 170 will not work in the new Models 145 and 175.

IBM 6611 Model 145 is suitable for building a backbone in a location with a number of connections. It can handle 8 LANs for 16 serial connections.

IBM 6611 Model 175 is the largest 6611 model, which provides seven adapter slots that can support the connection of a maximum of 14 LAN ports or 28 WAN ports or a combination of LAN and WAN ports, each at less than their maximum capacity. Thus, IBM 6611 Model 175 is a solution for large regional headquarters and campuses.

2.3.1.4 New Adapters

The new adapter features apply to all models of the 6611. These adapters include a new processor and twice the memory of the previous 6611 adapters. The following are the benefits of the new adapters:

- **Increased port density**

 New LAN adapters, which offer either two token-ring or Ethernet ports, are now available; a new WAN adapter is added which provides four serial ports. This doubles the number of LAN and WAN ports previously available for the 6611.

- **LAN/WAN combinations**

 In addition, two new combination adapters are introduced, each offering one LAN port (either token-ring or Ethernet) plus two WAN serial ports on a single adapter. This allows maximum flexibility while preserving adapter slots in all models.

- **Improved performance**

 In general, the new adapters perform better than the old adapters. A four-port serial adapter can fully load four serial lines at T1 speeds. At E1 speeds, the four-port serial adapter performs better than two of the old two-port serial adapters. A token-ring serial combination adapter can handle all of the traffic that previously/program could be handled by two adapters (a token-ring and a two-port serial adapter). In the case of an Ethernet serial combination adapter, if the serial interfaces are heavily used with small frame sizes, there is a slight reduction on the Ethernet maximum throughput due to the processing power being shared with the serial interfaces.

 The 6611PERF package on MKTTOOLS provides in-depth information on performance. Your IBM account representative will be able to provide you with a copy of this document.

- **Increased connectivity**

 All new adapters with multi-interface serial ports, including the new combination adapters, can support any of the following physical interfaces on any port, including a mix of different interfaces per card:

 – CCITT V.35 - at speeds from 9600 bps to 2.048 Mbps

 – CCITT V.36 - at speeds from 9600 bps to 2.048 Mbps

 – EIA 422/449 - at speeds from 9600 bps to 2.048 Mbps

 – EIA 232/CCITT V.24 - at speeds from 4800 bps to 19.2 kbps

 – CCITT X.21 - at speeds from 4800 bps to 2.048 Mbps

 Selection of the interface is determined by the adapter cable. So, if a change in the network interface equipment is required in the future, only a new cable is needed to switch interfaces.

- **Investment protection**

 These adapters are all supported on Models 140 and 170 as well as the new models. This enables customers with installed 6611s to exploit the versatility and performance improvements of these new adapters without requiring an investment in a new platform.

The following is a list of all of the types of adapters which will be available for any 6611 Model (note that Model 120 is available only in fixed configurations).

Different adapters must be ordered depending on whether you're putting the adapters in a Model 125 or a Model 145/175. The new adapter types are:

- Four-port multi-interface serial adapter
- Two-port token-ring network 16/4 adapter
- Two-port Ethernet adapter
- Multi-interface serial/token-ring combination adapter
- Multi-interface serial/Ethernet combination adapter
- Two-port multi-interface serial adapter (new, reduced cost)
- One-port token-ring network 16/4 adapter (new, reduced cost)
- One-port Ethernet adapter (new, reduced cost)

The existing adapters are:

- Four-port SDLC adapter
- X.25 adapter

Note

The four-port SDLC adapter and the X.25 adapter are unchanged. The new processor and double the memory used by the new adapters are not applicable to the four-port SDLC and X.25 adapters.

2.3.2 Multiprotocol Connectivity

The IBM 6611 Network Processor provides routing of the network layer protocols used by the following protocol suites:

- Internet Protocol (IP)
- Novell NetWare Internetwork Packet Exchange (IPX)
- Xerox Network Systems (XNS) Internet Transport Protocol
- DECnet Phase IV and DECnet Phase IV-Prime
- AppleTalk Phase 2
- Banyan VIrtual NEtworking Systems (VINES)

2.3.2.1 Communication Adapter Features Supported

The communication adapter features supported for each of the protocols that can be routed by the IBM 6611 Network Processor are summarized in Table 12.

Table 12 (Page 1 of 2). IBM 6611 Adapter Ports and Supported Protocols

Adapter Ports	Ethernet			Token-Ring		Serial			SDLC	X.25
Standard	Version 2	IEEE 802.3		IEEE 802.5						CCITT X.25
Framing / Protocols	Type	LLC	SNAP	LLC	SNAP	PPP	Frame Relay	Token-Ring Bridge Prgm	SDLC	X.25
IP	X		X		X	X	X	X		X
XNS	X	X	X	X	X	X	X	X		
IPX•	X	X	X	X	X	X	X	X		X

Table 12 (Page 2 of 2). IBM 6611 Adapter Ports and Supported Protocols

Adapter Ports	Ethernet			Token-Ring		Serial			SDLC	X.25
Standard	**Version 2**	**IEEE 802.3**		**IEEE 802.5**						**CCITT X.25**
Framing / Protocols	**Type**	**LLC**	**SNAP**	**LLC**	**SNAP**	**PPP**	**Frame Relay**	**Token-Ring Bridge Prgm**	**SDLC**	**X.25**
AppleTalk			X		X	X	X	X		
DECnet	X				X	X	X	X		
Banyan VINES	X		X	X	X	X	X	X		
SNA•	X		X	X		X	X	X	X	X
APPN•	X		X	X		X	X	X		X
NetBIOS•	X		X	X		X	X	X		X
Source-route Bridging				X	X	X	X	X		
Transparent Bridging	X	X	X			X	X			
Translational Bridging	X	X	X	X	X	X	X	X		

Note:

- Also supports native Novell 802.3 for IPX.
- To run APPN, DLSw must be configured. APPN also requires that DLSw or IP be configured for APPN network nodes to communicate across a WAN.
- For local DLSw of SNA, the configuration of IP is not required. For remote DLSw of SNA and NetBIOS, IP must be configured on the link between DLSw session partners.

All of the protocol suites that are supported for a communication adapter feature can be used concurrently across the same communication adapter interface. For example, an interface on the Multi-Interface Serial Adapter can be configured to support the transport of TCP/IP, NetWare, XNS, DECnet and AppleTalk protocol suites concurrently.

This is possible because the data link protocols used by the communication adapter features that support multiple protocol suites provide a mechanism for distinguishing between the various protocol suites sharing the same communication interface.

For example, the PPP data link protocol uses a 2-byte protocol code within each frame to distinguish between protocol suites sharing the same communication interface.

Note: The communication adapter features supported for the TCP/IP protocol suite can also be used to support the transfer of information that originates from nodes that use either the SNA or the NetBIOS protocol suites. This is achieved using the IBM 6611 Network Processor data link switching function which encapsulates the SNA or NetBIOS protocols inside the TCP protocol. This is described further in topic 2.3.4, “Data Link Switching” on page 145.

2.3.2.2 Routing Table Maintenance

The IBM 6611 Network Processor uses separate routing tables for each of the protocol suites it supports. That is, there is one routing table for each protocol suite supported by the IBM 6611 Network Processor.

For the DECnet, XNS, NetWare, AppleTalk and Banyan VINES protocol suites, their routing tables are maintained using the corresponding routing table maintenance protocol dynamically. For example, the XNS protocol suite uses XNS RIP (Routing Information Protocol) for this purpose.

For the TCP/IP protocol suite, several routing table maintenance protocols can be used either singularly or in combination to maintain the single TCP/IP routing table. Additionally, static routes can be manually defined during configuration of the IBM 6611 Network Processor.

The TCP/IP routing table maintenance protocols supported by the IBM 6611 Network Processor are:

- Interior protocols used within an autonomous system:

 TCP/IP RIP (Routing Information Protocol)

 Hello

 OSPF (OSPF)

- Exterior protocols used between autonomous systems:

 EGP (Exterior Gateway Protocol)

 BGP (Border Gateway Protocol)

2.3.2.3 Filtering

The IBM 6611 Network Processor multiprotocol routing function provides a very comprehensive filtering capability. There are three types of filtering provided:

1. Filtering based on protocol suite

 The routing of each supported protocol suite can be selectively disabled or enabled for each IBM 6611 Network Processor. That is, each IBM 6611 Network Processor can be configured to either ignore (filter) or route each of the supported protocol suites.

 For example, an IBM 6611 Network Processor can be configured to ignore the token-ring segments DECnet protocol suite, and only route the TCP/IP, XNS, AppleTalk and NetWare protocol suites. Frames received by the IBM 6611 Network Processor that are identified as DECnet will be discarded, and frames received that are identified as either TCP/IP, XNS, AppleTalk or NetWare will be routed.

2. Filtering based on communication interface

 If the routing of a particular protocol suite is enabled for an IBM 6611 Network Processor, it can be selectively disabled or enabled for each communication interface. That is, each communication interface can be configured to either ignore or route a particular protocol suite.

 For example, an IBM 6611 Network Processor that is enabled for routing the TCP/IP protocol suite, can be configured to ignore the TCP/IP protocol suite on one of its communication interfaces, and only route the TCP/IP protocol suite on the remaining communication interfaces.

3. Filtering based on network layer address

For each protocol suite the IBM 6611 Network Processor provides additional filtering capabilities that allow the enabling or disabling of routing based on network layer addresses. These filters are either specific to a particular communication interface or global to all communication interfaces.

The specifics of these filters vary between protocol suites as each protocol suite uses a different form of network layer addressing.

2.3.3 Bridging with IBM 6611

The 6611 supports routing and three types of bridging:

- Source-route bridging

 Source-route bridging is used on the 6611 to bridge frames between token-ring LANs.

- Transparent bridging

 Transparent bridging is used on the 6611 to bridge frames between Ethernet LANs.

- Translational bridging

 Translational bridging allows you to bridge frames between token-ring and Ethernet LANs.

The following topics provide a brief description of bridging with 6611.

2.3.3.1 Source-Route Bridging

Source-route bridging is used to interconnect networks at the data link layer of the OSI reference model. Source-route bridging involves forwarding MAC frames based on information in the MAC header. A frame is passed from bridge to bridge until it reaches the final destination.

A bridge examines each frame to determine whether it is destined for the bridge itself or for another device. The bridge uses data from its tables or information in the frame header to determine whether the frame should be forwarded to another device. Source-route bridging depends on the device that sends the frame (the source) to indicate, within the frame, the complete route to the final destination. The route is a sequence of identifiers for the bridges and rings along the path from the source to the destination device.

Unlike a router, a bridge does not examine the network protocol header that is imbedded in the data field of the MAC frame. The bridge is unaware of the network protocol information in the data field. Consequently, a bridge is sometimes referred to as protocol independent.

The 6611 can be configured to provide local or remote bridge functions.

Local Bridge Function: A single 6611 can be used to interconnect multiple token-rings that are directly attached to the 6611. Figure 84 on page 133 illustrates this local bridge function.

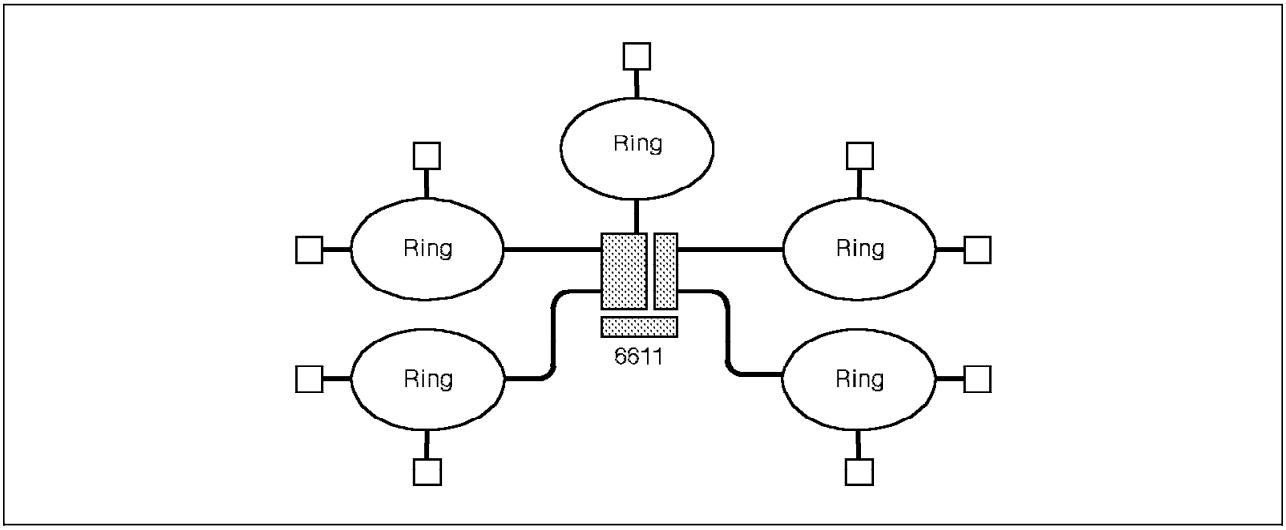


Figure 84. Local Source-Route Bridge Function

Each token-ring segment is attached to the IBM 6611 Network Processor using an IBM 6611 Token-Ring Network 16/4 Adapter. IBM 6611 Network Processor can be used to interconnect two or more token-ring segments across an intervening frame relay network or telecommunication link.

The IBM 6611 Network Processor when used as a source-route bridge can forward three types of frames:

All-Routes Broadcast: When the IBM 6611 Network Processor receives an all-routes broadcast frame on one of its token-ring interfaces, it copies the frame to all the other IBM Token-Ring Network segments to which it is attached. In doing so it updates the RI (Routing Information) field of each copy of the received frame with its bridge number, and the segment number of the destination token-ring segment. The RI field is also updated with the source segment number if it is not already present within the RI field.

Single-Route Broadcast: When the IBM 6611 Network Processor receives a single-route broadcast frame, it only copies the frame to the other token-ring segments if the corresponding interface has been enabled for the forwarding of single-route broadcast frames. Each interface can either be manually or automatically configured for the forwarding of single-route broadcast frames. The RI field for each copy of the received frame is updated in the same manner as for all-routes broadcast frames.

Non-Broadcast with Routing Information Field: When the IBM 6611 Network Processor receives a non-broadcast frame that contains an RI field it will forward the frame if the next entry in the RI field contains the bridge number of the IBM 6611 Network Processor and the segment number of a segment attached to the IBM 6611 Network Processor.

The IBM 6611 Network Processor is able to participate in the automatic configuration of the single-route broadcast function using the spanning tree algorithm with other source-route bridges that support this capability.

Remote Bridge Function Between 6611s: Two 6611s can be used to interconnect two or more token-rings across an intervening frame relay network or telecommunications link. Figure 85 on page 134 shows two sample

configurations that use this remote bridge function. The function is sometimes called native mode bridging, to distinguish it from the remote bridge function described below.

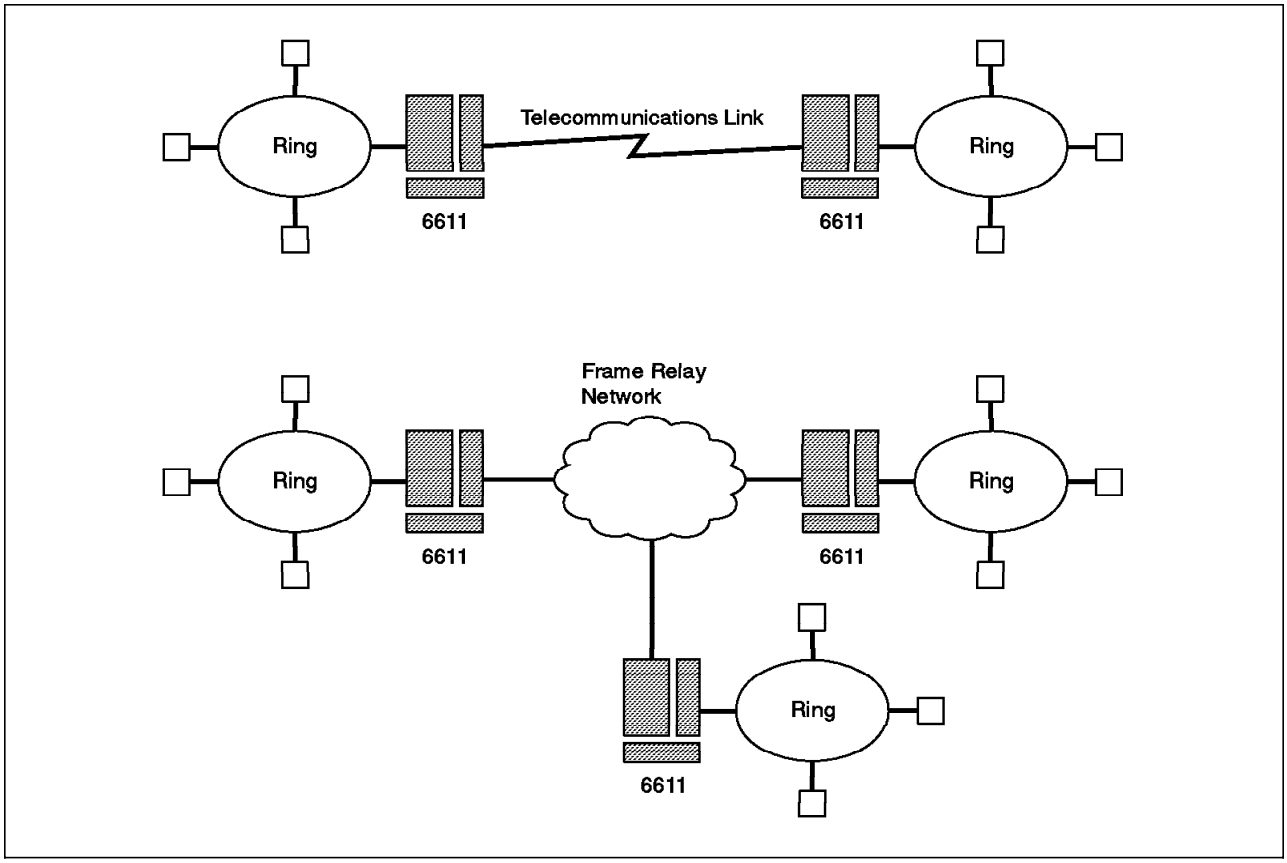


Figure 85. Remote Source-Route Bridge between 6611s

Each token-ring segment is attached to an IBM 6611 Network Processor using an IBM 6611 Token-Ring Network 16/4 Adapter. The remote connections between each IBM 6611 Network Processor can utilize the two multi-interface serial ports, and can use either the PPP or frame relay data link protocols.

Each connection between IBM 6611 Network Processors can be either:

- A point-to-point communication facility such as the T1 or E1 services provided by many common carriers. Such a connection would use PPP data link protocols.
- A DLC (Data Link Connection) across a frame relay service. Many DLCs can share the same physical interface to a frame relay service using a unique DLCI (Data Link Connection Identifier) to distinguish between each DLC. This allows an IBM 6611 Network Processor to establish connections with many other IBM 6611 Network Processors using a single physical interface to a frame relay service.

The bridge number assigned to the IBM 6611 Network Processor will be used not only for bridging with remote token-ring segments attached to other IBM 6611 Network Processors, but also for local bridging and remote bridging with PS/2s.

Remote Bridge Function Between a 6611 and a PS/2: The IBM 6611 supports remote bridging between a 6611 and a PS/2 workstation running either the IBM

Token-Ring Network Bridge Program, Version 2.2, or the IBM Remote Token-Ring Bridge/DOS, Version 1.0.

Figure 86 shows a sample configuration using this remote bridge function. The function is sometimes called compatibility mode bridging. In this configuration, the 6611 functions as the primary half of the bridge and the Bridge Program functions as the secondary half of the bridge. A telecommunications link connects the 6611 to the PS/2 workstation running the bridge program. The devices communicate using a proprietary protocol.

Note

The proprietary protocol used on the telecommunications link is referred to as the LAN Bridging Protocol within the 6611 library.

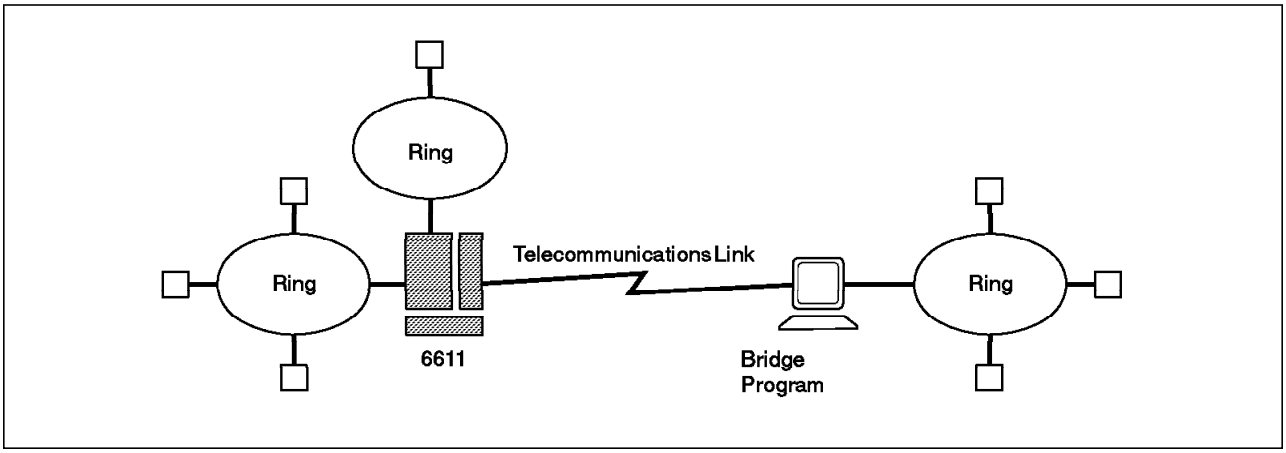


Figure 86. Remote Source-Route Bridge between a 6611 and a PS/2 Workstation Running a Bridge Program

Token-ring segments are attached to the IBM 6611 Network Processor using the IBM 6611 Token-Ring Network 16/4 Adapter. Remote connections between IBM 6611 Network Processors and PS/2s utilize point-to-point protocol (PPP), and can be attached to the IBM 6611 Network Processor using the two multi-interface serial ports.

The bridge number assigned to the IBM 6611 Network Processor will be used not only for bridging with remote token-ring segments attached via PS/2s, but also for local bridging and remote bridging with other IBM 6611 Network Processors.

Additionally, one of the token-ring segments locally attached to the IBM 6611 Network Processor must be selected to become the designated ring. All of the PS/2 remote bridges connected to an IBM 6611 Network Processor are logically bridged to the designated segment. An example of how to use a designated ring is shown on Figure 87 on page 136.

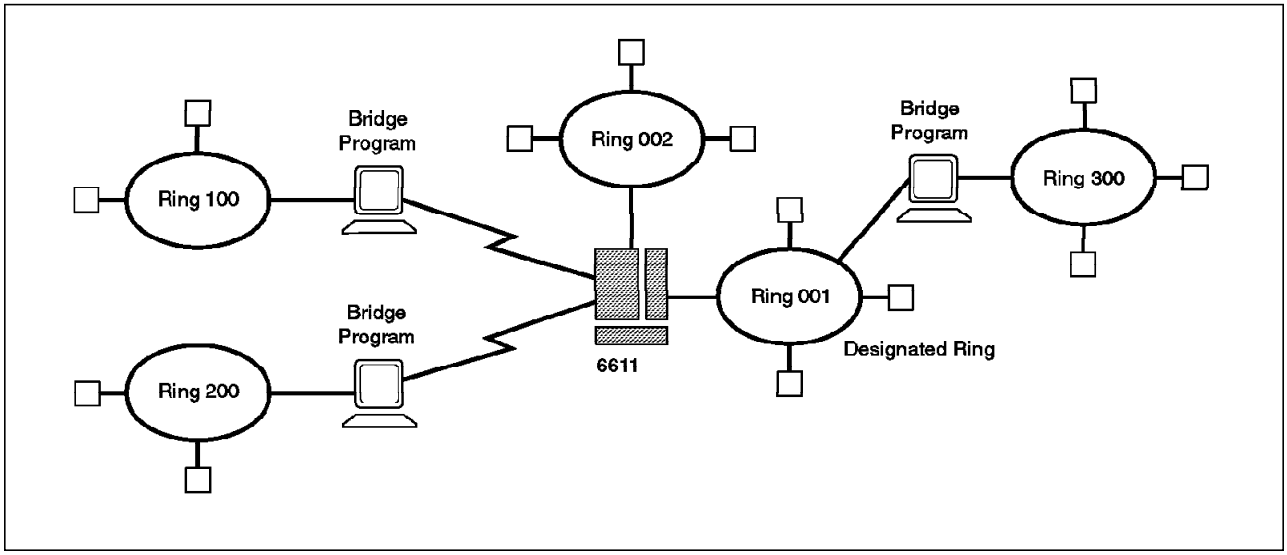


Figure 87. Remote Source-Route Bridge and the Designated Ring

Note: Frames transported by the IBM 6611 Network Processor between token-ring segments other than the designated segment do not appear on the designated segment. Instead they are processed entirely within the IBM 6611 Network Processor. However, the designated segment number does appear in the RI field of frames transported to or from remote token-ring segments attached to PS/2 remote bridges.

Filtering: The IBM 6611 Network Processor source-route bridging function provides a very comprehensive filtering capability.

Filters can be configured for each communication interface that participates in source-route bridging. This includes interfaces on both the IBM 6611 Token-Ring Network 16/4 Adapter and the Multi-Interface Serial Adapter when remote source-route bridging is used.

For each communication adapter interface, both inbound and outbound filters can be configured. Inbound filters act upon frames received by the IBM 6611 Network Processor across the communication interface. Outbound filters act upon frames scheduled for transmission by the IBM 6611 Network Processor across the communication interface.

There are five types of filters which can be configured for each interface. With the exception of the hop count filter, each type can be configured separately for inbound and outbound operation. The five filter types available are:

Hop Count: This filter can be used to process frames that have more than an allowable number of hops in their RI (Routing Information) field.

MAC Address: This filter can be used to process frames that are to or from specific MAC (media access control) addresses.

Source SAP: This filter can be used to process frames that contain a specific source SAP (service access point).

SNAP Value: This filter can be used to process frames that contain a specific SNAP header. SNAP headers exist in frames that have source and destination SAP values of X′AA′.

Segment Number: This filter can be used to process frames that contain a specific origin segment number within the RI (Routing Information) field.

Each type of filter only acts upon either single-route broadcast, or all-routes broadcast frames, or both. Each type of filter can be set to operate in one of two modes:

- Include only frames which match the filter characteristic (not used by the hop count filter). This is *permit* mode.
- Exclude only frames which match the filter characteristic (always used by the hop count filter). This is *deny* mode.

With the exception of the hop count filter, each type of filter provides the capability for multiple values to be filtered concurrently, and a mask capability allows a range of values to be specified with a single entry. Only those bits set in the mask are used for comparisons between the value specified and the frame being processed by the filter.

All five types of filters can be used concurrently if required. With the exception of the hop count filter, each type of filter can be individually enabled or disabled.

Notes

Use of the SNAP value filter requires that the corresponding source SAP filter also be enabled. For example, to use the outbound SNAP value filter for an interface, the outbound source SAP filter for the same interface must also be enabled. No SAPs need be defined for the source SAP filter if only the SNAP value filter is required.

The hop count filter can be effectively disabled by setting the hop count value to 7 (seven) which is the maximum hop count possible in token-rings.

To illustrate how multiple filters work together, consider the following scenario where outbound source SAP, outbound ring number and hop count filters are used concurrently for a token-ring interface. The filter settings are listed in Table 13.

Table 13. Example Filter Settings

Filter Type	Mode	Value(s)
Outbound Source SAP	Deny	X′AA′ X′F0′
Outbound Ring Number	Permit	X′100′ X′200′ X′300′
Hop Count	Deny	2

For a frame to pass through the interface for which these filters are enabled, it must meet all of the following criteria:

1. It must have a source SAP that is not X′AA′ or X′F0′ (as indicated by the filter settings in the list). For example, a frame with a source SAP of X′04′ would pass this filter, but a frame with a source SAP of X′F0′ would not.
2. It must contain an origin segment number of X′100′, X′200′ or X′300′. For example, a frame with a routing information field of X′100 1 300 0′ would meet this requirement, whereas a frame with a routing information field of X′400 1 300 0′ would not.

3. The routing information field must contain two hops or less. For example, a frame with a routing information field of X′100 1 200 1 300 0′ would meet this requirement, whereas a frame with a routing information field of X′200 1 800 1 100 1 300 0′ would not.

2.3.3.2 Transparent Bridging

Transparent bridging, like source-route bridging, is a method used to interconnect networks at the data link layer. The 6611 supports Ethernet transparent bridging, as defined in the IEEE standard for Media Access Control Bridges (802.1D).

In source-route bridging, the device sending a frame discovers the preferred route to a destination device and that route is included within the frame transmitted by the sending device. In transparent bridging, a sending device transmits frames without regard for the location of a destination device. The bridges in the network are responsible for forwarding each frame to its proper destination.

Transparent bridges receive all frames transmitted on the LAN segments to which they are attached, and examine the source and destination addresses of each frame. By examining the source address of a frame, the bridge learns the port and LAN segment associated with a sending device. This information is stored in a routing table or filtering database and is used to make future decisions about how to forward frames. By examining the destination address of a frame arriving on a port, the bridge determines if the frame should be forwarded to another port or discarded (the destination device and sending device, in this case, are on the same side of the bridge). Each adapter maintains its own filtering database.

Transparent bridges, like source-route bridges, do not examine the network protocol header imbedded in the data field of the MAC frame. The bridge is unaware of the network layer protocols and bridges all frames independently of these protocols.

The 6611 can be configured to provide the following transparent bridge functions:

- Local bridging
- Remote bridging

Local Bridging Function: A single 6611 can be used to interconnect multiple Ethernet LANs that are directly attached to the 6611. Figure 88 on page 139 illustrates this local bridging function.

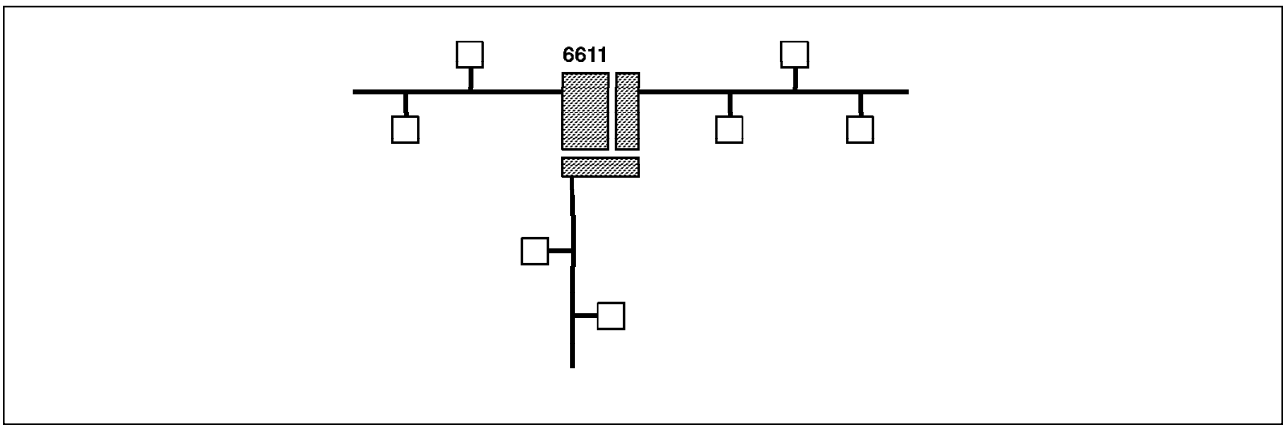


Figure 88. Local Transparent Bridge Function

Remote Bridging Function: Two 6611s can be used to interconnect two or more Ethernet LANs across an intervening frame relay network or telecommunications link. Figure 89 on page 140 shows several 6611 configurations using the remote bridging function. As indicated in the figure, Ethernet and token-ring frames can be transported over the same telecommunications link or frame relay connection.

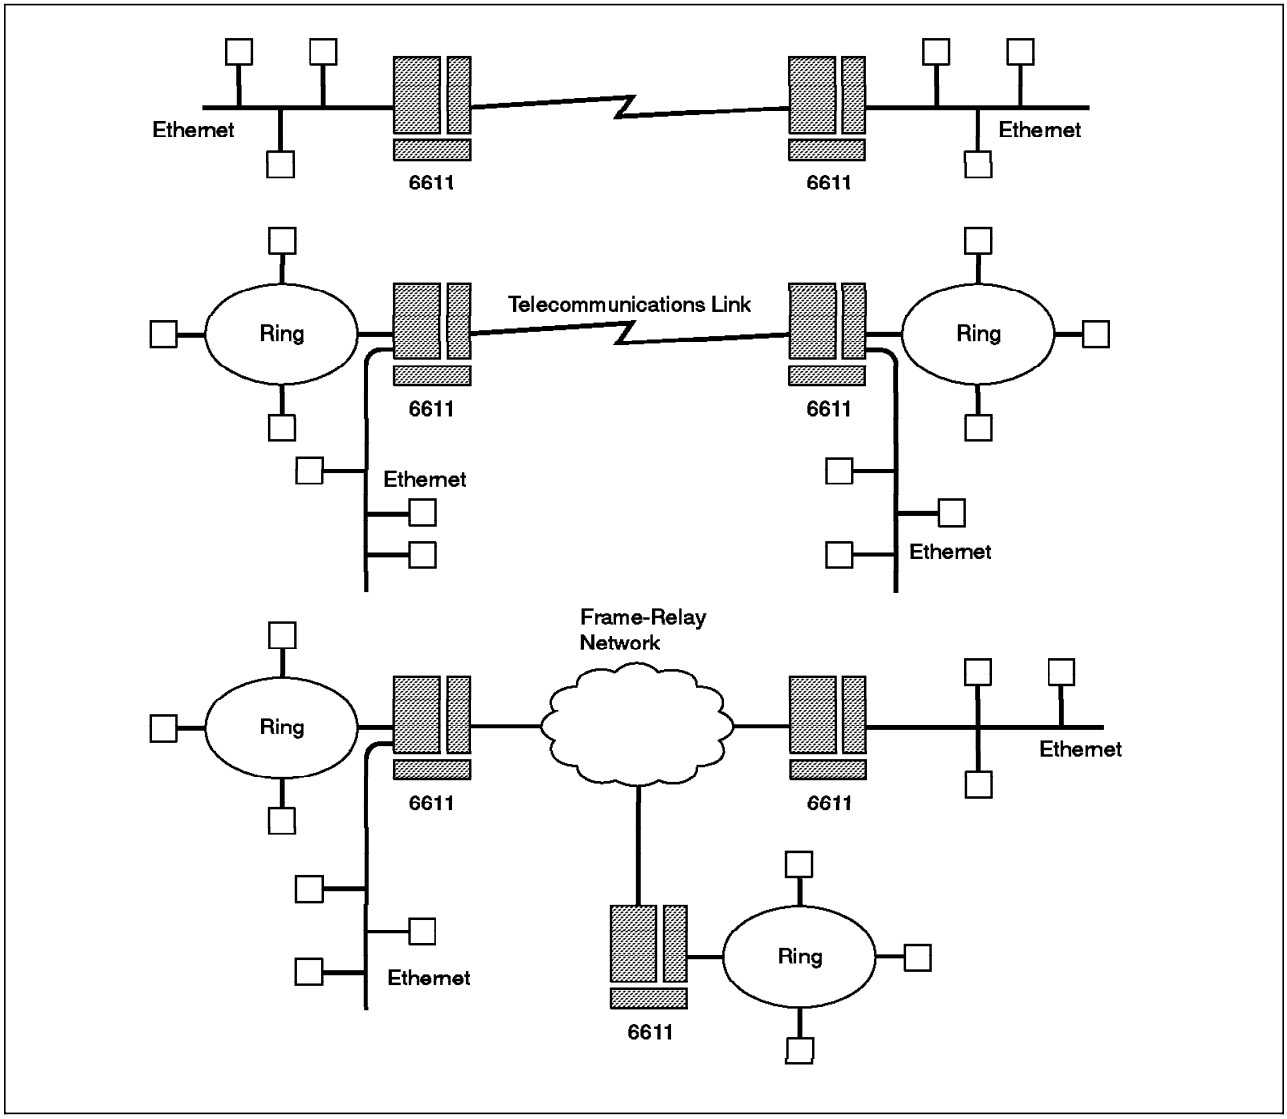


Figure 89. Remote Transparent Bridge Function

2.3.3.3 Translational Bridging

On the 6611, token-ring ports can be configured to support source-route bridging, and Ethernet ports can be configured to support transparent bridging. Because each LAN type uses a different frame format and bridging technique, token-ring and Ethernet LANs cannot be interconnected without providing a method of translation. *Translational bridging* is the method used on the 6611 to bridge frames between these different LAN types. Translational bridging, as implemented on the 6611, is sometimes referred to as *source-route transparent bridging* (SRTB or SR-TB).

When you configure the 6611 node as a translational bridge, it operates in the following manner:

- If the source and destination ports for a frame use the same bridging technique, the frame is bridged between the ports without translation.
- If the source and destination ports for a frame use different bridging techniques, the translational bridge converts the frame into the format required for the destination LAN, and bridges the frame.

Frames in IEEE 802.5 format (for token-ring LANs) will be converted to either Ethernet Version 2.0 or IEEE 802.3 format as required by the destination Ethernet LAN. Ethernet frames will be converted to IEEE 802.5 format as required.

To a device on a token-ring LAN, the 6611 translational bridge appears as a source-route bridge. To a device on an Ethernet LAN, the translational bridge is functionally transparent. To enable it to interconnect token-ring and Ethernet LANs, the translational bridge maintains two address databases, as follows:

- The Ethernet database contains the source addresses for stations detected on Ethernet LANs and the frame format that each station uses for data transmission (Ethernet V2.0 or IEEE 802.3).
- The token-ring database contains the source addresses and routing information for stations on token-ring LANs that have forwarded frames to Ethernet LANs.

Notes

- The translational bridging function on the 6611 is compatible with functions provided by the IBM 8209 and 8229 LAN Bridge products.
- The 6611 does not support source-routing transparent (SRT) bridging, which combines source-route bridging and transparent bridging techniques into a single bridging method for token-ring LANs.

The 6611 translational bridge can be configured to provide the following bridge functions:

- Local bridge function
- Remote bridge function between two 6611 translational bridges
- Remote bridge function between a 6611 translational bridge and a 6611 source-route bridge or transparent bridge
- Remote bridge function between a 6611 translational bridge and a PS/2 workstation running either the IBM Token-Ring Network Bridge Program Version 2.2, or IBM Token-Ring Network Bridge/DOS Version 1.0

Local Bridging Function: A single 6611 can interconnect multiple token-ring and Ethernet LANs that are directly attached to the 6611. Figure 90 on page 142 illustrates this local bridge function.

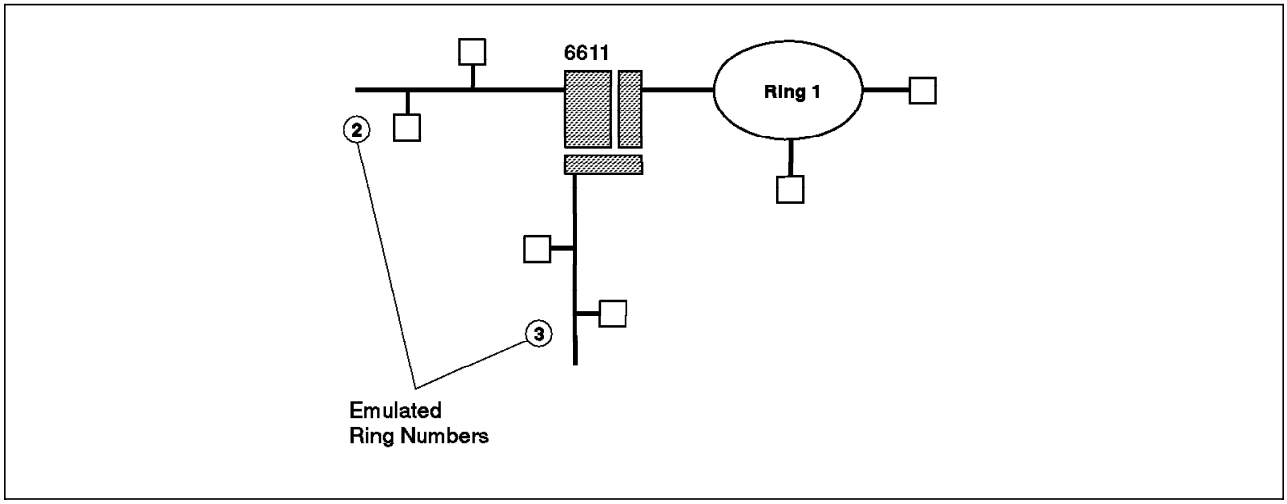


Figure 90. Local Translational Bridge Function

Remote Bridging Function Between 6611s: Two 6611s can be used to interconnect token-ring and Ethernet LANs across an intervening frame relay network or telecommunications link. Figure 91 on page 143 shows two sample configurations that use this remote bridge function. The recommended method for connecting two 6611 translational bridges is to configure *dual mode bridging* on each end of the serial link. When you configure dual mode bridging, bridged frames are translated only if the source and destination LANs require different MAC frame formats.

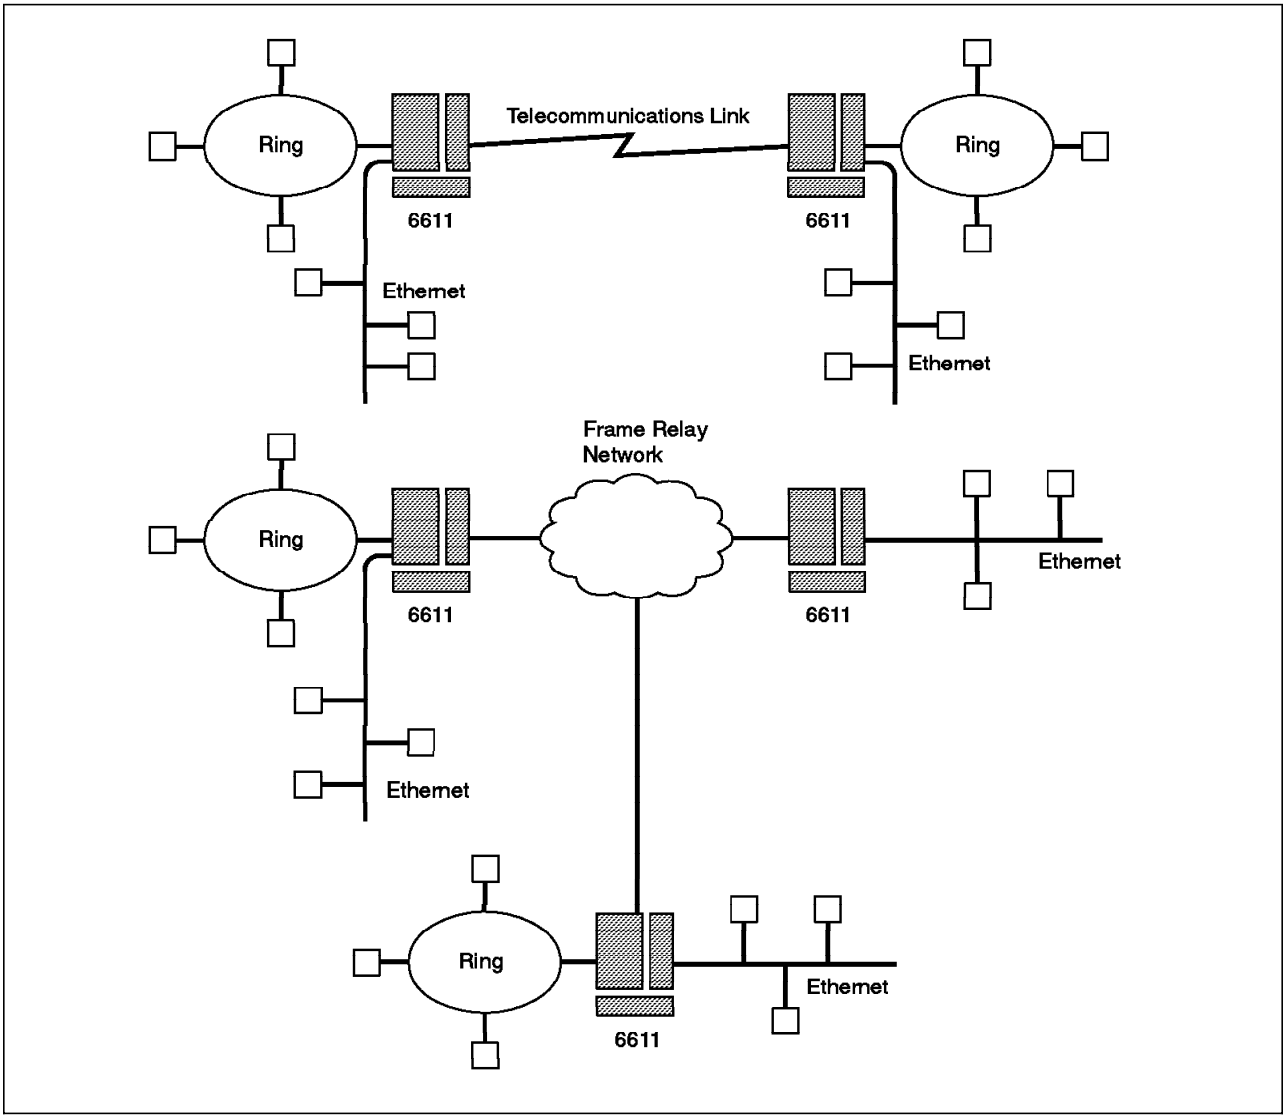


Figure 91. Remote Bridging Function between 6611 Translational Bridges

Remote Bridging Function between a 6611 Translational and Non-Translational Bridge: A 6611 translational bridge can be connected to a 6611 source-route bridge or transparent bridge across an intervening frame relay network or telecommunications link. The LANs attached to each bridge can communicate across the WAN connection. Figure 92 on page 144 shows a sample configuration that uses this remote bridging function.

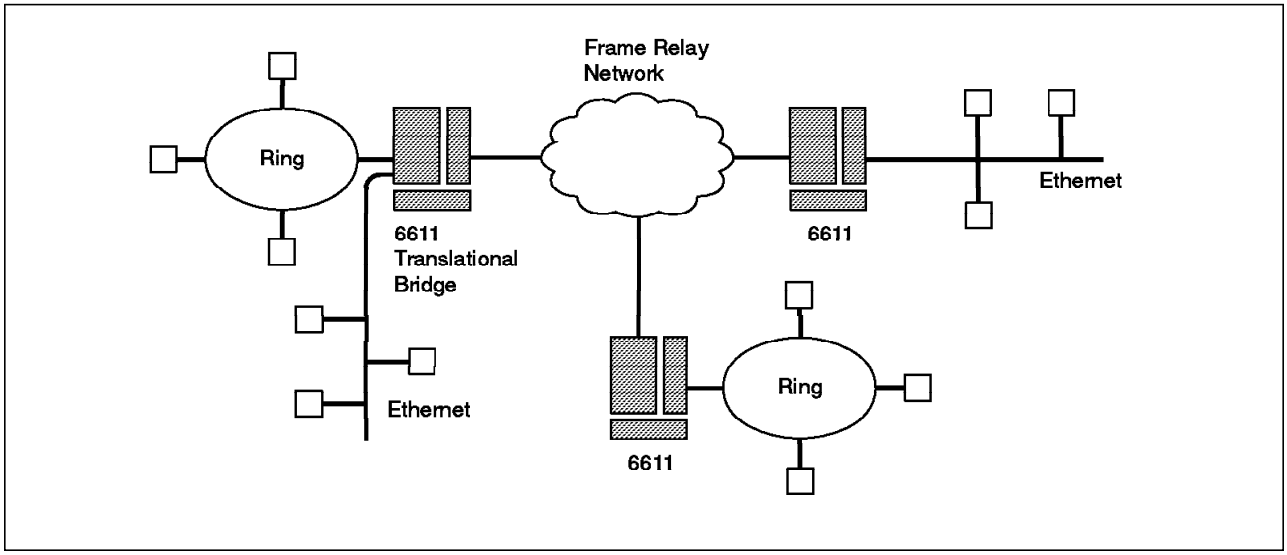


Figure 92. Remote Bridging Function between a Translational and a Non-Translational Bridge

Remote Bridging Function between a 6611 Translational and a PS/2: On remote bridging between a 6611 translational bridge and a PS/2 workstation running either the IBM Token-Ring Network Bridge Program Version 2.2, or the IBM Remote Token-Ring Bridge/DOS Version 1.0, the frames can be bridged between 6611 ports configured for source-route, transparent, or dual mode bridging and the PS/2 workstation running the bridge program.

Figure 93 on page 145 shows a sample configuration using this remote bridging function. The function is sometimes called *compatibility mode bridging*. In this configuration, the 6611 functions as the primary half of the bridge, and the bridge program functions as the secondary half of the bridge. A telecommunications link connects the 6611 to the PS/2 workstation running the bridge program. The devices communicate using a proprietary protocol.

Note

The proprietary protocol used on the telecommunications link is referred to as the *LAN Bridging Protocol* within the 6611 library.

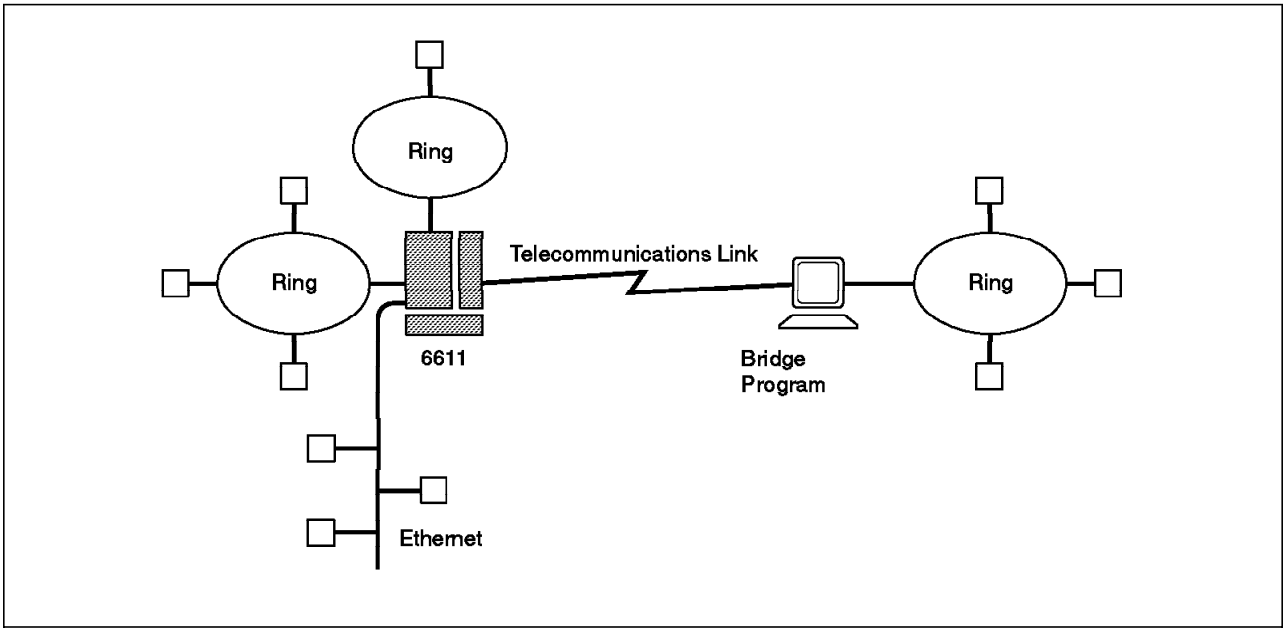


Figure 93. Remote Bridging Function between a Translational Bridge and a PS/2 Workstation Running a Bridge Program

2.3.3.4 Coexistence with Other IBM Bridge Products

The IBM 6611 Network Processor can coexist with other bridges, such as the IBM 8209 or IBM 8229 and the IBM Personal System/2, using the IBM Token-Ring Network Bridge Program Version 2.2. This includes support for automatic single-route broadcast configuration using the spanning tree algorithm.

However, the IBM 6611 Network Processor does not implement the following functions provided by other IBM bridge products:

- RPS (Ring Parameter Server)
- REM (Ring Error Monitor)
- CRS (Configuration Report Server)
- LRM (LAN Reporting Mechanism)
- LBS (LAN Bridge Server)

As a consequence, there are some limitations when using IBM LAN Network Manager to manage interconnected token-rings that incorporate IBM 6611 Network Processor-based bridges.

2.3.4 Data Link Switching

DLSw is a method of transporting SNA and NetBIOS frames.

The DLS function provides the capability to integrate the transport of the NetBIOS and SNA protocol suites with the other protocol suites that can be routed by the IBM 6611 Network Processor.

Devices that make use of the DLS function are configured as if they were directly attached to each other via a single data link or data link network.

In reality these devices only have a direct data link or data link network connection to an IBM 6611 Network Processor. The IBM 6611 Network Processor then transports information received on the data link or data link network connection to another IBM 6611 Network Processor. This second IBM 6611 Network Processor has a direct data link or data link network connection with the ultimate destination device.

The two data links or data link networks that are connected via the DLS function need not be the same type of data link or data link network. For example, an SNA device attached via an SDLC data link to a 6611 Network Processor can use the DLS function to connect to an SNA device attached via a token-ring network data link network.

The DLS function uses the TCP transport layer protocol (part of the TCP/IP protocol suite) to implement a transport network between IBM 6611 Network Processors. This transport network can comprise many intermediate nodes, data links and data link networks, if required, through the use of the IP network layer protocol (also part of the TCP/IP protocol suite).

Note

Intermediate nodes in the transport network used to connect IBM 6611 Network Processors that are providing the DLS function do not have to be IBM 6611 Network Processors, provided that they can support the IP network layer protocol.

A TCP connection is automatically established between each pair of IBM 6611 Network Processors that are participating in the DLS function across the TCP/IP transport network. To support the establishment of these TCP connections, each IBM 6611 Network Processor is configured with the TCP/IP network addresses of the other IBM 6611 Network Processors participating in the DLS function.

It is possible to configure an IBM 6611 Network Processor to accept incoming DLS TCP connections from other IBM 6611 Network Processors without explicitly configuring the other IBM 6611 Network Processors. This may reduce the amount of configuration effort required to set up complex DLS environments. However, at least one of the two IBM 6611 Network Processors participating in each DLS TCP connection must be configured with the TCP/IP network address of the other IBM 6611 Network Processor.

The communication adapter features that can be used with the DLS function fall into the following four categories:

- Those that support direct data links to SNA devices
- Those that support direct data links to NetBIOS devices
- Those that support indirect data links to token-ring devices (both SNA and NetBIOS) via a remote source-route bridge configuration
- Those that support connection to the TCP/IP transport network used to interconnect IBM 6611 Network Processors that provide the DLS function

The DLS function incorporates several features to reduce the need to send data across the TCP/IP network that interconnects the IBM 6611 Network Processors participating in the DLS function.

The key feature is the cache in which each IBM 6611 Network Processor maintains a table of remote SNA and NetBIOS devices along with the IBM 6611 Network Processor that is able to reach that remote device through the fastest path. Each IBM 6611 Network Processor constructs its cache dynamically by sending queries to other IBM 6611 Network Processors only when needed. The cache can be preloaded with default entries when the IBM 6611 Network Processor is configured to further reduce the need for queries to be sent to other IBM 6611 Network Processors.

An age out timer is used to remove old cache entries after a period of time. The timeout used by the age out timer can be set when the IBM 6611 Network Processor is configured.

Note

At the time of writing, the cache used by the DLS function could only be used to locate the MAC addresses of remote SNA and NetBIOS devices. As a consequence, NetBIOS requests to locate particular NetBIOS names were copied to all interfaces enabled for DLS on all IBM 6611 Network Processors that participate in the DLS function. However, it is intended that the cache be used to locate NetBIOS names of remote NetBIOS devices. This would dramatically reduce the number of NetBIOS broadcasts that flow across the TCP/IP network that interconnects all IBM 6611 Network Processors participating in the DLS function.

To explain how data link switching is implemented in the 6611, we define two types of data link switching: *local data link switching* and *remote data link switching*. In local data link switching, the data link switching function is performed within a single 6611. In remote data link switching, stations attached to two or more 6611s communicate across an IP network using data link switching. The following topics summarize the features of the two types of data link switching.

There are several differences in the operation of the DLS function for SNA and NetBIOS devices. For this reason each is described separately in 2.3.4.3, “SNA Data Link Switching” on page 151 and in 2.3.4.4, “NetBIOS Data Link Switching” on page 153.

For more information about DLSw networking considerations, see Chapter 4 of *Local Area Network Concepts and Products: LAN Architecture*, SG24-4573.

2.3.4.1 Local Data Link Switching

Local data link switching is used for SNA transport only. It supports communication between a LAN-attached SNA device and a synchronous data link control (SDLC) secondary station that is link-attached to the 6611. The LAN-attached SNA device may be on a LAN directly attached to the 6611, or it may be on a remote LAN that is joined to the 6611 by one or more bridges.

The SDLC secondary station must be a physical unit (PU) type 2.0 or 2.1 and must be operating in normal response mode. During configuration of the 6611, the secondary station is assigned a MAC sub-layer address so that it appears to other network devices to be on a LAN.

Local data link switching converts SDLC frames to IEEE 802.2 LLC type 2 frames. Bridging is used to transport the converted frames (SNA frames encapsulated in

a MAC sub-layer frame) to a directly attached LAN or to the next bridge in the path of an interconnected LAN. The local data link switching function does not convert token-ring MAC sub-layer frames to Ethernet MAC sub-layer frames. However, a route to an interconnected LAN may contain a bridge, such as an IBM 8209 or 8229 LAN Bridge, that converts token-ring MAC sub-layer frames to Ethernet MAC sub-layer frames. A technique called spoofing is used to send acknowledgments to the source station from the 6611 to which the source station is attached, instead of from the destination station.

When configuring local DLSw, configuration of DLSw partners and IP routing is optional.

A sample local data link switched network is shown in Figure 94 and in Table 14 on page 149.

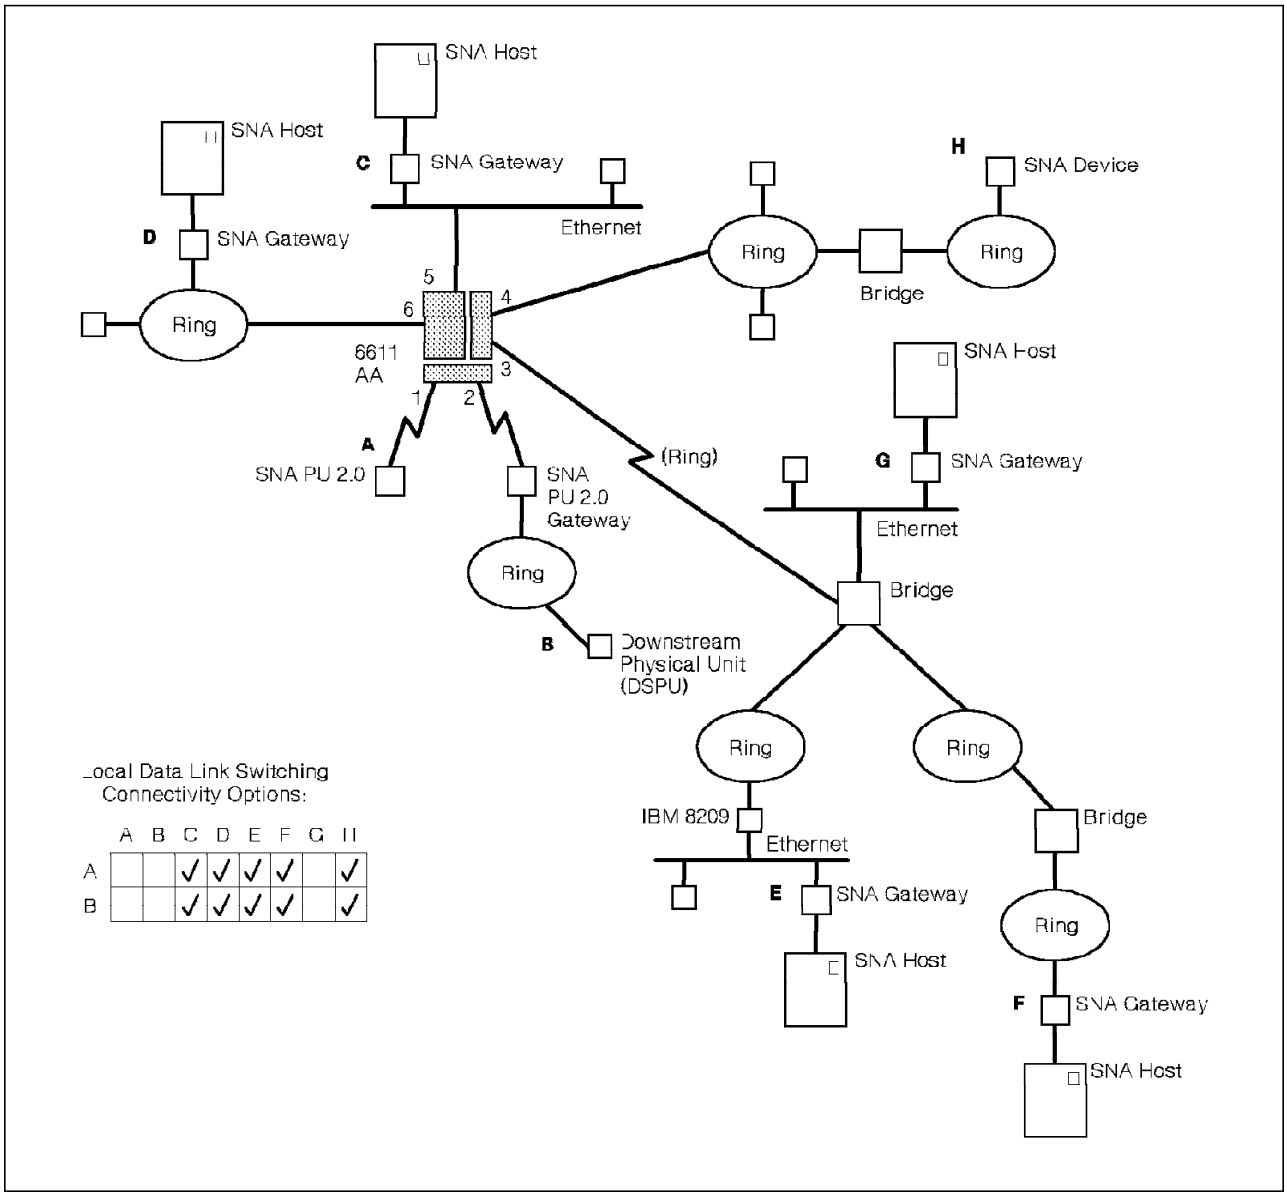


Figure 94. Sample Local Data Link Switched Network

Table 14. Sample Local Data Link Switched Network		
Reference	**Configuration Item**	**Node-Level or Port-Level Configuration**
AA	6611	Source-route bridging, transparent bridging, DLSw for SNA
1	SDLC port	SDLC, SNA
2	SDLC port	SDLC, SNA
3	Serial port	Source-route bridging, DLSw for SNA
4	Token-ring port	Token-ring, source-route bridging, DLSw for SNA
5	Ethernet port	Ethernet, transparent bridging, DLSw for SNA
6	Token-ring port	Token-ring, source-route bridging, DLSw for SNA

2.3.4.2 Remote Data Link Switching

Remote data link switching is used for both SNA and NetBIOS transport. An SNA or NetBIOS station attached to a 6611 uses remote data link switching to communicate with an SNA or NetBIOS station attached to another 6611. SNA stations may be link-attached or LAN-attached to the 6611s; NetBIOS stations must be LAN-attached. The 6611s, called partners, must be configured for data link switching. The partners communicate with each other across an IP network.

- SDLC-to-LAN communication across a WAN

 Remote data link switching performs SDLC-to-IEEE 802.2 type 2 conversion. This permits a link-attached SDLC secondary station to communicate with a LAN-attached SNA device.

- LAN-to-LAN communication across a WAN

 Remote data link switching supports communication between SNA or NetBIOS stations on token-rings and Ethernets. Remote data link switching can convert token-ring MAC sub-layer frames to Ethernet MAC sub-layer frames, and conversely, so that devices on token-rings and Ethernets can communicate with each other.

The 6611s communicate with the SNA and NetBIOS stations using IEEE 802.2 LLC type 2. The LLC connections are terminated at the 6611s. Spoofing is used to send acknowledgments to the source station from the 6611 to which the source station is attached, instead of from the destination station. This reduces traffic on the WAN.

The hop count for source-route bridging is also terminated at the 6611s. Thus, the source station may be up to 7 hops from the first 6611 in the path and the receiving station may be up to 7 hops from the last 6611 in the path.

For transport between the data link switching partners, the SNA or NetBIOS frames are encapsulated in IP datagrams. The partners communicate with each other using TCP. The route between two partners can contain IP routers that are not 6611s, as long as they are compatible with the 6611. The 6611s in an IP route between partners must be configured for IP routing, but they need not be configured for data link switching.

A sample remote data link switched network is shown in Figure 95 on page 150.

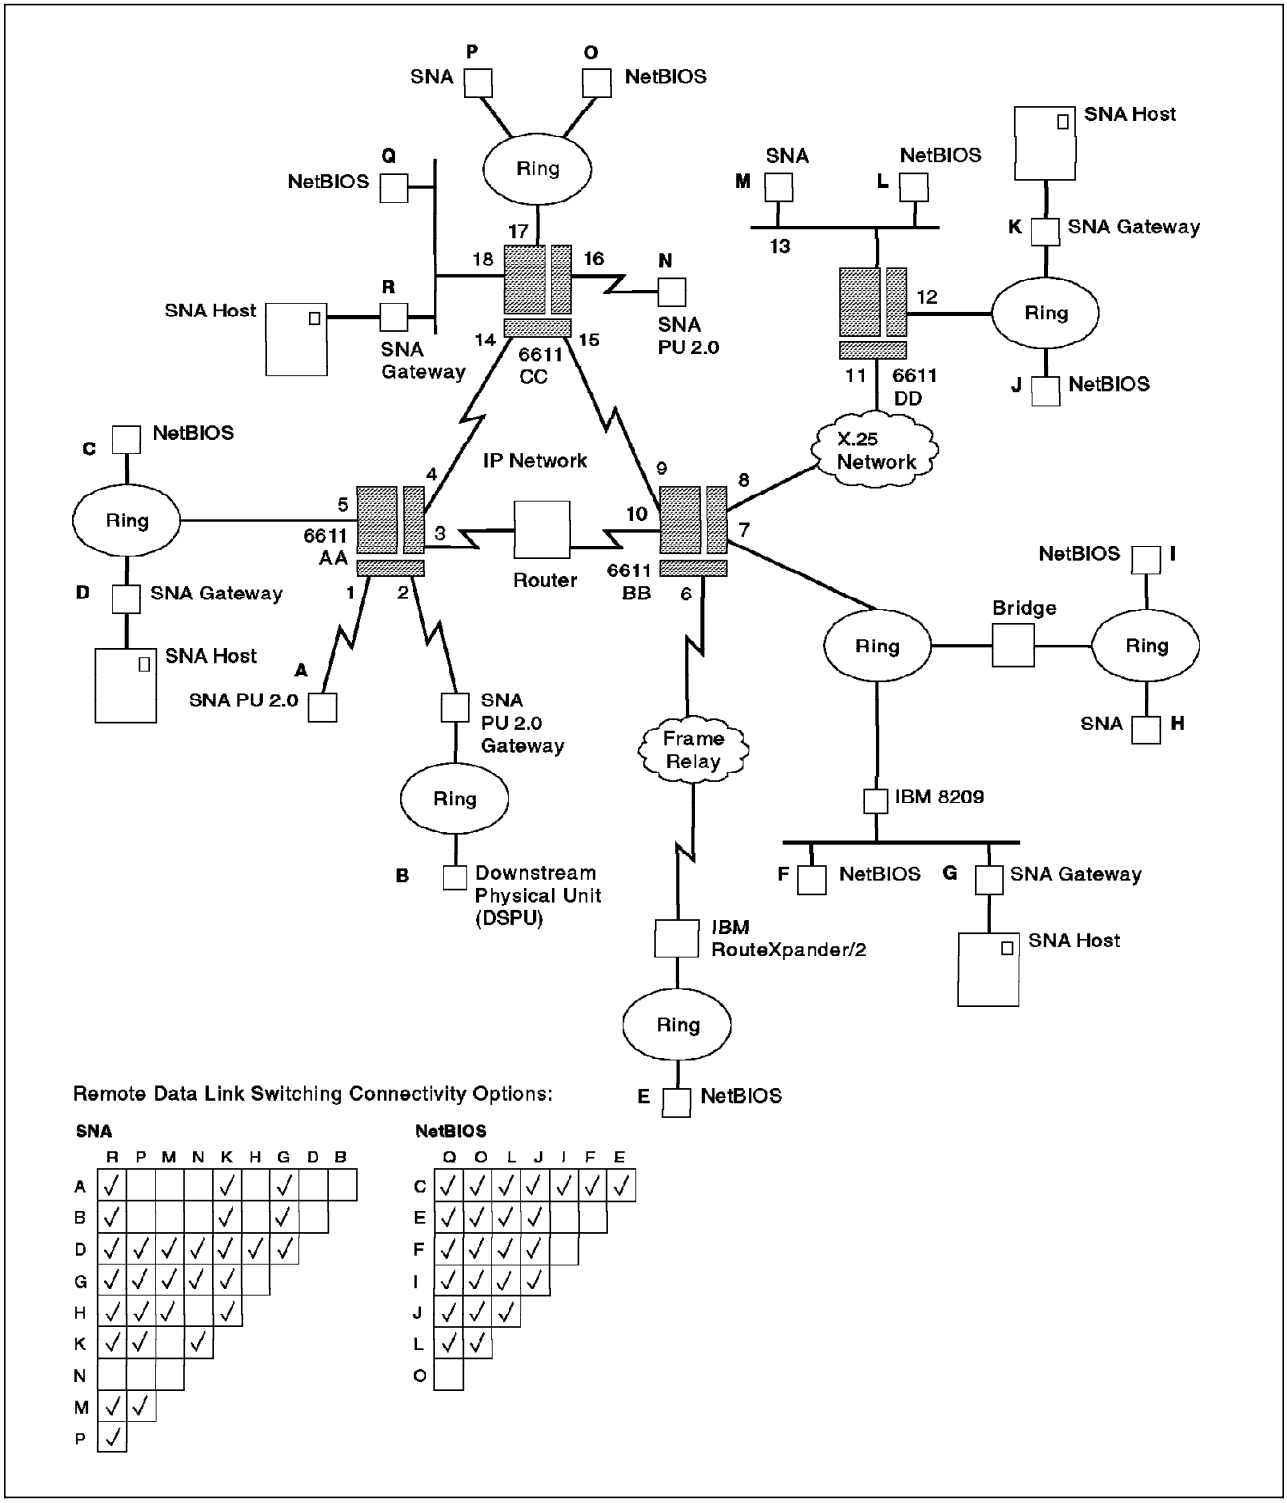


Figure 95. Sample Remote Data Link Switched Network

The node-level and port-level configurations for the 6611s in Figure 95 are summarized in Table 15 on page 151.

Table 15. Configuration of the Sample Remote Data Link Switched Network		
Reference	**Configuration Item**	**Node-Level or Port-Level Configuration**
AA	6611	OSPF, source-route bridging, DLSw for SNA and NetBIOS
BB	6611	OSPF, source-route bridging, IP over X.25, DLSw for SNA and NetBIOS
CC	6611	OSPF, source-route bridging, transparent bridging, DLSw for SNA and NetBIOS
DD	6611	OSPF, source-route bridging, IP over X.25, DLSw for SNA and NetBIOS
1	SDLC port	SDLC, SNA
2	SDLC port	SDLC, SNA
3	Serial port	PPP, IP
4	Serial port	PPP, IP
5	Token-ring port	Token-ring, source-route bridging, DLSw for SNA and NetBIOS
6	Serial port	Frame relay, source-route bridging, DLSw for SNA and NetBIOS
7	Token-ring port	Token-ring, source-route bridging, DLSw for SNA and NetBIOS
8	X.25 port	X.25, IP
9	Serial port	PPP, IP
10	Serial port	PPP, IP
11	X.25 port	X.25, IP
12	Token-ring port	Token-ring, source-route bridging, DLSw for SNA and NetBIOS
13	Ethernet port	Ethernet, transparent bridging, DLSw for SNA and NetBIOS
14	Serial port	PPP, IP
15	Serial port	PPP, IP
16	SDLC port	SDLC, SNA
17	Token-ring port	Token-ring, source-route bridging, DLSw for SNA and NetBIOS
18	Ethernet port	Ethernet, transparent bridging, DLSw for SNA and NetBIOS

2.3.4.3 SNA Data Link Switching

The DLS function supports the interconnection of SNA devices attached to either a token-ring or an SDLC multipoint non-switched line. A typical example of the use of the DLS function for SNA devices is illustrated in Figure 94 on page 148 and in 2.3.4.1, “Local Data Link Switching” on page 147.

As a prerequisite for the DLS function, each participating token-ring segments IBM 6611 Network Processor that supports token-ring-attached SNA devices, must be configured to support source-route local bridging on all token-ring interfaces used with the DLS function.

Note

Local bridging will be used in preference to the DLS function to provide connections between token-ring-attached SNA devices that are connected to the same IBM 6611 Network Processor via different token-ring segments.

Each IBM 6611 Network Processor participating in the DLS function must also be configured with a virtual segment number. This virtual segment number must be the same for all IBM 6611 Network Processors participating in the DLS function.

Additionally, SNA devices attached to an IBM 6611 Network Processor via an SDLC multipoint non-switched line are assigned a token-ring LAA (locally administered address), SAP (Service Access Point) and SNA XID (Exchange ID). These will be used by the IBM 6611 Network Processor to represent such devices to other SNA devices that are using the DLS function.

Note

A single hop is used in the RI (Routing Information) field to reach an SNA device accessible via the DLS function from a token-ring segment directly attached to a IBM 6611 Network Processor. Therefore, SNA devices can be, at most, six hops from an IBM 6611 Network Processor to reach SNA devices accessible via the DLS function.

The DLS function only supports the attachment of SNA devices via SDLC multipoint lines that are of PU (Physical Unit) Type 2.0. The attachment of PU Type 2.1 devices is not supported unless they provide a PU 2.0 compatibility mode. The attachment of PU Type 4 devices (such as the IBM 3745 Communications Controller) is not supported either.

There are two consequences of this:

1. SDLC-attached devices cannot establish connections with other SDLC-attached devices. This is because SNA PU type 2.0 devices cannot directly communicate with each other as peers.
2. SDLC-attached devices can only support a single connection to another SNA device attached to a token-ring. The other SNA device will usually be a PU type 4, such as the IBM 3745, or a PU type 5.

DLSw SNA Traffic Prioritization: This function was implemented in the Multiprotocol Network Program Version 1 Release 3 (MPNP). It can be defined as a method that allows SNA frames to have adequate priority over NetBIOS frames. It applies to the DLSw traffic from all the ports on the 6611. Additional priority can be given to SNA frames by a two-pronged approach as follows:

1. **SNA/NetBIOS Ratio (Bias)**

 The user can specify the ratio of how many SNA frames are to be sent per NetBIOS frame. Valid SNA/NetBIOS ratio settings are from 0 to 9. If the ratio is set at 9, nine SNA frames will be transmitted on the link per NetBIOS frame. The frames are selected from the DLSw data stream preserving the order of the frames.

 There is no capability that allows NetBIOS frames to have priority over SNA frames. This function is for increasing the priority for SNA traffic.

2. **NetBIOS Frame Size Reduction**

 NetBIOS tends to send frames as large as the transport mechanism will allow, while SNA tends to send very small frames. This can often lead to NetBIOS using most of the transport's bandwidth. The NetBIOS largest frame size option allows users to force the frames to be broken into segments. In other words, NetBIOS will be forced to use smaller frames, thus allowing SNA Bias to have a more predictable effect. The choices of the valid largest allowed NetBIOS frame in bytes are 2052, 1500 and 516.

2.3.4.4 NetBIOS Data Link Switching

The DLS function supports the interconnection of NetBIOS devices attached to either a token-ring or a CSMA/CD (Carrier Sense Multiple Access/Collision Detection) LAN using either DIX Ethernet V2 or IEEE 802.3 frame formats. A typical example of the DLS function for NetBIOS devices is illustrated in Figure 95 on page 150.

NetBIOS devices on token-rings are handled in a similar way to SNA devices on token-rings. That is, remote NetBIOS devices will appear as if they are on the DLS virtual segment.

NetBIOS devices on CSMA/CD LANs cannot be handled in a similar way to that used for SNA devices on token-rings. Instead, the ability of NetBIOS to dynamically bind a MAC address to a NetBIOS name is exploited.

From the perspective of NetBIOS devices on CSMA/CD LANs, all remote NetBIOS devices appear as if they have the MAC address of the 6611 Ethernet Adapter. This is possible because the NetBIOS protocol discovers the MAC address of other NetBIOS devices using broadcast frames sent to the NetBIOS functional address.

2.3.4.5 Estimating DLSw Storage Requirements

Developing a DLSw configuration requires careful design and planning for efficient utilization of available system resources. To assist you in planning your configuration and determining your 6611 memory needs, IBM provides a storage estimating tool called the IBM 6611 Storage Estimate EXEC. For information on this tool, contact your IBM marketing representative and ask for the Internetworking Marketing Specialist for your trading area.

Memory expansion features are available if additional memory is required for the 6611. An 8 MB memory expansion (feature code 4008) is available on Models 125, 145, and 175. A 16 MB memory expansion (feature code 4016) is available on Models 145 and 175. The 16 MB memory expansion for Models 140 and 170 is available by RPQ 8Q1414.

2.3.5 IBM 6611 Network Processor Enhancements - Release 4

There are many enhancements that will be available on IBM 6611 - Release 4 that we can emphasize:

- High Performance Routing (HPR), with the following features:
 - Automatic Network Routing (ANR) is a sophisticated new source-routing method that delivers unmatched price/performance for mission-critical data.
 - Rapid Transport Protocol (RTP) allows safe reroute of data around failed links or notes.

 - Adaptive Rate-Based (ARB) provides superior flow and congestion control.
- Dependent LU Requester (DLUR) which enables dynamic configuration of dependent LUs.
- Enhanced Priority Queueing Support with three new HPR data queues enrich the 6611′s priority queueing scheme.
- FR Boundary Access Node (BAN) that provides the ability to bridge token-ring and Ethernet SNA traffic directly to an FEP (3745) without frame conversion by DLSw router.
- Frame Relay RFC 1490 is a standard that specifies how SNA and multiprotocol LAN traffic can be natively and efficiently encapsulated in frame relay frames for transport across a wide-area network.
- ITU-T LMI Support via Frame Relay - ITU-T Q.9333 Annex is a standard that defines means of status and the notification of outage for frame relay PVC.
- DLSw V1 Compliance RFC 1795 is an industry-standard method for transmitting SNA and NetBIOS traffic across a TCP/IP wide area network.
- Support for RFC 1027; Transparent Subnetting which enables the 6611 to act as a transparent subnet ARP gateway.
- Support for RFC 1542; BOOTP which enables the 6611 to act as an BOOTP relay agent. Also allows the 6611 to act as a relay agent for host RFC 1534.
- 2210 EasyStart that allows the IBM 6611 Network Processor to act as a BOOTP relay agent for 2210s which needs to download its initial configuration information from the network.
- IPX Filtering enhancements with new IPX RIP filters that allows a network administrator to filter inbound and outbound RIP filters using network numbers ranges; one filter can be applied to all ports.
- Fast IPL Time for the 6611 Network Processor has been significantly improved.
- Auxiliary Power Shutdown restricts shutdown of UPS.
- System Manager Enhancements where several new enhancements to the System Manager function are provided.
- Up to 32 MB of Memory Upgrade for M125 enable Customers to order additional memory (up to 32 MB) for their 6611 Model 125 using Feature Code 4008.
- DASD Size Enhancement where new models of 6611 will begin shipping with larger hard drives.
- OS/2 & DOS/WIN Configuration Transfer Support for sending configurations through the network using TCP/IP socket connection to a 6611.
- Multiple Retrieve Function that provides the ability to retrieve configurations files from multiple routers for configuration updates.

Chapter 3. Additional IBM Software Solution

IBM offers complete end-to-end Internet solutions, so that customers can get Internet enabled at every point from initial access to creating an Internet presence, integrating the Internet into core business applications and enabling true networked applications. These capabilities leverage offerings from virtually every corner of IBM, including Lotus. This chapter covers the Internet offerings, including TCP/IP, Internet Connection software products and Lotus InterNotes. This chapter does not include any discussion of hardware platforms. It is IBM′s intention to enable all platforms, including Intel, AIX, PowerPC, AS/400 and S/390, for the Internet.

For the most current information on IBM′s Internet offerings, see the IBM Internet home page at the URL http://www.ibm.com/internet/ and the Lotus home page at the URL http://www.lotus.com.

3.1 Overview

IBM offers a set of products and services that help customers get connected to the Internet quickly, easily and securely. These offerings support systems ranging from desktop and laptop computers to UNIX workstations and PS/2s, and from AS/400 business computers to the S/390 mainframe.

IBM′s offerings span hardware (which is not covered here), software (both for the client and server) and network and consulting services (which are not covered here).

Software

This includes client software for accessing and browsing the Web and server software for Web information management, gateway services, firewall, and Web authoring and application building tools. Some Lotus software is also covered in this chapter.

TCP/IP client/server software

- IBM TCP/IP Version 2 Release 3 for VM
- IBM TCP/IP Version 3 Release 1 for MVS
- IBM TCP/IP Version 2.1.1 for DOS
- IBM TCP/IP Version 3.0 for OS/2
- IBM AIX for RISC System/6000 Version 4.1.4 (TCP/IP included)
- IBM OS/400 Version 3 Release 2 (TCP/IP included)

Client software

- Internet Connection for OS/2 Warp
- Internet Connection for Windows
- Warp Connect
- WebExplorer for AIX
- Secure WebExplorer for AIX
- Secure WebExplorer for OS/2 Warp

© Copyright IBM Corp. 1996

Server software

- IBM Internet Connection Server for OS/2 Warp and AIX
- IBM Internet Connection Secure Server for OS/2 Warp and AIX
- IBM Internet Connection Server for MVS/ESA
- IBM Internet Connection Secure Server for MVS/ESA
- IBM WebConnection for OS/400

Internet servers

- IBM Internet POWERsolution for AIX-IBM Internet Connection
- IBM Internet POWERsolution for AIX-Netscape

Lotus InterNotes

- Lotus InterNotes Web Publisher
- Lotus InterNotes News

Firewall software

- IBM Internet Connection Secured Network Gateway for AIX

Information Gateways

- IBM DB2/WWW
- IBM CICS/WWW
- IBM MQ Series/WWW

World Wide Web Tools

- IBM VisualAge WWW
- IBM Electronic Publishing Edition
- IBM Hyperwise

Network Services

These are dial or leased-line connections to the Internet and network applications. (It is the application code that actually runs on IBM Global Network backbone and is sold as a subscription service.)

IBM Global Network Internet Connection

- Dial
- Leased line
- Firewall service

IBM Global Network Content Services

- Hosting
- Design and creation

IBM InfoMarket Service

This is the first secure environment for intellectual property owners to reach a world-wide audience over the Internet. InfoMarket acts as a clearinghouse for commercial content and service providers giving them greater control over distribution.

For further information about InfoMarket Service, refer to Chapter 12, “Networked Applications” on page 523.

Consulting Services

These are professional services to assist clients in planning, designing and implementing Internet solutions. This includes Web site design and development, business and information technology consulting, I/T security solutions, installation services and education.

- Business Transformation Services
- I/T Consulting-Internet Consulting Services
- Internet Planning and Design Workshops
- Internet Implementation ISO
- Interactive Media Design (Advanced Internet Graphics and Design)
- Internet Connection SNG Firewall Installation
- AS/400 Gopher Client Installation
- WebConnection for OS/400 Smoothstart Installation
- Internet Connection Server Smoothstart Installation
- Customer Seminars and Education
- I/T Security Consulting and Services

For further information about IBM Consulting, refer to Chapter 14, “Consulting Services” on page 553.

Note

Firewall software, Information Gateways, World Wide Web Tools and Networking and Consulting Services are covered by other chapters in this book.

3.2 TCP/IP Client/Server Software

Table 16 (Page 1 of 2). Operating Systems and Their Corresponding TCP/IP Applications

	S/370			PC		RISC/6000	AS/400
	MVS	VM	AIX	DOS	OS/2	AIX	OS/400
FTP	c/s	c/s	c/s	c/s	c/s	c/s	c/s
TELNET	c/s	c/s	c/s	c/	c/s	c/s	c/s
TN3270	c/s	c/s	c/	c/	c/	c/	c/s
SMTP	c/s	c/s	c/s	c/s	c/s	c/s	c/s
SUN RPC	c/s	c/s	c/s	c/	c/s	c/s	c/s
NFS V2	/s	/s	c/s4	c/	c/s	c/s4	c/s
NCS	c/s	c/s			c/s	c/s	
X Window	c/	c/	c/		c/s	c/s	
REXEC	c/s	c/s	c/s	c/	c/s	c/s	
TFTP		c/	c/s	c/s	c/s	c/s	
LPR/LPD	c/s	c/s	c/s	c/s	c/s	c/s	c/s

Table 16 (Page 2 of 2). Operating Systems and Their Corresponding TCP/IP Applications

	S/370			PC		RISC/6000	AS/400
	MVS	**VM**	**AIX**	**DOS**	**OS/2**	**AIX**	**OS/400**
SNMP	m/a	m/a		m/a	m*/a	m/a	/a
Sockets	c/s	c/s	c/s	c/s	c/s	c/s	c/s
Kerberos	c/s	c/s			c/s		
DNS	r/s	r/s	r/s	r/	r/s	r/s	r/
TALK			c/s		c/s	c/s	
Finger			c/s	c/	c/	c/s	
PING	x	x	x	x	x	x	x
NETSTAT	x	x	x	x	x	x	x
RIP	x	x	x	x	x	x	

Note:

4 = support SUN PC-NFS 4.0

c/s = client/server support

m/a = monitor/agent support, monitor for DOS: NetView for Windows

r/s = resolver/server support

x = noted function exists for the product

Further information about TCP/IP can be found in the *TCP/IP Tutorial and Technical Overview*, GG24-3376-04.

3.3 Client Software

The following sections refer to the IBM client software offerings.

3.3.1 Internet Connection for OS/2 Warp and Windows

Internet Connection for OS/2, included in OS/2 Warp, OS/2 Warp Connect and OS/2 Warp Connect 4.0 Beta (Merlin), and Internet Connection for Windows are easy-to-use tools that provide quick and easy access to the Internet. These products lets you electronically subscribe to IGN Internet Connection Services or choose another Internet Service provider that supports the serial line Internet protocol (SLIP) or point-to-point protocol (PPP) methods of communication.

In addition to the World Wide Web, Internet Connection gives you access to other popular Internet applications and functions:

- E-mail
- Gopher
- News Reader
- Viewer
- Archie
- Basic TCP/IP functions

IBM WebExplorer is our browser. It provides an easy to use and interactive graphical user interface to the WWW. WebExplorer for OS/2 Warp is included in Internet Connection for OS/2 Warp. The browser included in Internet Connection for Windows is WebExplorer Mosaic, which is code that we licensed from Spyglass.

Enhancements to WebExplorer for OS/2 Warp were announced in September 1995 and include the following:

- Mail-to support. When selected, an easy-to-use form will come up that allows you to enter a message that is then mailed to the recipient specified in the mail-to tag.
- News articles are now displayed as a hierarchical tree making it easier to follow the thread of articles. Users can also easily post and subscribe to news groups.
- WebExplorer has improved integration with the Workplace Shell. You can drag a Uniform Resource Locator (URL) from the WebExplorer application to create a URL Workplace Shell object. This object can then be dropped back onto the WebExplorer application or just onto the WebExplorer icon triggering it to access the URL. A user can effectively organize, sort and categorize their favorite Web locations by using Workplace Shell folders.
- WebExplorer also supports document streaming. All supported image formats will be displayed using the streaming method, thus improving the performance and presentation of the images.

3.3.2 Warp Connect

Warp Connect includes the same code and functions as Internet Connection for OS/2 Warp with the addition of a LAN connection.

3.3.3 Secure WebExplorer for OS/2 Warp and AIX

In addition to all the features of the base WebExplorer product, this supports Secure Hypertext Transfer Protocol (S-HTTP) and Secure Sockets Layer (SSL). These technologies ensure that information is encrypted and arrives safely at its intended destination. Secure browsers and servers allow the user to conduct secure transactions on the Internet, such as online purchases using a credit card number.

3.3.4 WebExplorer for AIX

This code will be included in the AIX operating system. It includes the same functions as WebExplorer for OS/2 Warp.

3.4 Server Software

The following sections refer to the IBM server software offerings.

3.4.1 Internet Connection Server for OS/2 Warp and AIX

The IBM Internet Connection Server has the features needed to build home pages on the Internet. The IBM Internet Connection Server can:

- Act as a repository for home pages created with Hypertext Markup Language (HTML).
- Answer requests from a Web browser (client) using Hypertext Transfer Protocol (HTTP) to transfer documents.
- Provide proxy support, allowing a Web browser to access remote servers not directly accessible to it. The proxy server supports requests from HTTP, FTP, and Gopher and acts on their behalf.

- Support proxy caching by temporarily storing files and then quickly responding to the next request for the files.
- Provide application interfaces, using Common Gateway Interface (CGI); this is an emerging standard API between the Internet Connection Server and another application, such as a database.

An easy-to-use HTML form is provided to help you configure the IBM Internet Connection Server to meet your business needs. You can specify options such as time-out settings, proxy servers, and caching.

3.4.2 Internet Connection Secure Server for OS/2 Warp and AIX

In order to conduct commerce over the Internet, it is important to ensure that the transactions are secure. To provide maximum flexibility in secure environments, IBM′s Internet Connection Secure Servers for AIX and OS/2 Warp support the emerging standards, Secure HyperText Transfer Protocol (S-HTTP) and Secure Sockets Layer (SSL). These security technologies ensure that information is encrypted for privacy and arrives at its intended destinations.

These servers were designed to be quickly and easily configured using any industry-standard browser. The installation and configuration is menu-based and includes online help designed to assist an administrator with making the correct choice.

3.4.3 Internet Connection Server for MVS/ESA

Supporting the industry networking standards, Internet Connection Server for MVS/ESA can interoperate with other Internet servers and clients. The server:

- Acts as a repository for home pages created with HTML
- Serves requests from a Web browser using HTTP to transfer documents
- Acts as a repository for images, sound clips and video clips
- Enables direct access through a Common Gateway Interface (CGI) to existing applications and business data maintained by CICS, DB2 and IMS
- Uses MVS System Authorization Facility (SAF) to route authorization requests to an external security manager such as RACF to allow for increased protection for HTTP resources
- Provides proxy support
- Supports proxy caching
- Provides easy-to-use HTML form to configure the server
- Supports workstation users with Web browsers inside and outside the enterprise

3.4.4 Internet Connection Secure Server for MVS/ESA

In addition to providing all the features and functions of the base MVS server, this will incorporate security technologies to ensure that information is encrypted and arrives safely at its end destination. Secure browsers and servers allow the user to conduct secure transactions on the Internet, such as online purchases using a credit card.

3.4.5 WebConnection for OS/400

With WebConnection for OS/400, an AS/400 can become a repository and server of data for the Internet. Functions include:

- The HTTP server provides a mechanism where the AS/400 system can be the repository of server data for businesses on the World Wide Web. This allows business access across the Internet to potential customers via Web browsers such as IBM OS/2 Warp's WebExplorer. Local access is available on a LAN via TCP/IP. Available data includes audio, video images, portions of the database, and textual descriptions.
- Logging of World Wide Web Server access for tracking activity. This allows AS/400 owners to track who is accessing their servers and what parts are being accessed most often, giving feedback on levels of interest in products and services.
- Access to AS/400 applications via the Hypertext Markup Language (HTML) device driver. This is a key differentiator for OS/400. With this enhancement, applications developed natively on OS/400 may now use web browsers as clients for their applications. The Web browsers can be locally attached via TCP/IP or located anywhere in the world when attached via the Internet. This means that AS/400 users can develop Internet applications using their preferred native application development environment. With the HTML device driver, current OS/400 applications, except those using bidirectional character sets (BiDi) and Text Assist, are converted so that they may be displayed on a web browser. These applications can be enhanced so that in addition to text, they may incorporate graphics, image, audio, and video.
- Serial Line Internet Protocol (SLIP) asynchronous communication connections allow inexpensive, limited bandwidth access to the World Wide Web and Internet.
- Anonymous FTP support provides access to a selected portion of data on the AS/400 system that the public can access without a password or user identification.

3.4.6 IBM Connection Server Family

Features:

- Easy online configuration via standard HTML forms
- Optional remote configuration via HTML forms
- Proxy support to allow Web browsers to access remote servers
- Proxy caching to temporarily store files and respond to subsequent requests without delay
- Common Gateway Interface (CGI) support to add application intelligence behind your HTML forms
- Use of the two most popular security protocols:
 - Secure Sockets Layer (SSL)
 - Secure Hypertext Transfer Protocol (S-HTTP)
- IBM httpd API to extend the server's base functions
- Server-side allows you to dynamically insert information into an HTML document that the sever sends to a client
- Error message customization

- Multiple IP address support to keep multiple Web sites on a single Internet connection server
- Integration of DB2 and CICS Gateways support to access DB2 data and run CICS transaction, processing applications using standard Web browsers

The IBM Internet connection server family table shows some hardware and software requirements you need to be aware of when thinking of your Web server software installation/configuration for some of IBM′s platforms.

Table 17. IBM Web Server Hardware/Software Requirements

Operating Systems	Hardware/Software Requirements
AIX	• RISC/6000 or IBM Power Series Family. • 6 MB of free disk space to installing the server. • Additional 4 MB of free disk space for install both the DB2 and CICS gateway features. • AIX 4.1.3 or later. • Any communication hardware adapter supported by TCP/IP protocol stack to make network connections. • For the DB2 gateway: – DB2/6000. – 2.5 MB of free disk space in the usr/lpp partition. – 0.5 MB of free disk space in the root directory. • For the CICS gateway: – CICS/6000 2.1. – 1 MB of free disk space in the usr/lpp partition.
OS/2	• PS/2 or Personal Computer that can support OS/2 Warp 3.0. • 4 MB of free disk space to install the server. • Additional 7 MB of free disk space for install both the DB2 and CICS gateway features. • OS/2 Warp 3.0 or later, or OS/2 Warp Server. • A partition formatted using the High Performance File System (HPFS). • For the DB2 gateway: – DB2/2 1.2 or later. – 600 KB of free disk space; 2 MB is recommended when installing the sample DB2 Gateway application. • For the CICS gateway: – Access to a CICS for OS/2 Server. – CICS Client for OS/2 1.0 installed, including updates from Corrective Service Disk (CSD) 1. – 4.5 MB of free disk space.
Windows NT	• PS/2 or Personal Computer that can support Windows NT 3.51. • Approximately 4 MB of free disk space. • Microsoft Windows NT Server or Client 3.51 with TCP/IP configured. • A partition formatted using either the NT File System (NTFS) or the High Performance File System (HPFS). Use NTFS to get the file protections and permissions that it provides.
HP-UX	• An HP9000 Series 700 with HP-UX 10.01 or later with approximately 6 MB of free disk space to install the server. • HP-UX 10.01 or later.
Solaris	• A Sun SPARC station or UltraSPARC station. • Any communication hardware adapter that supports TCP/IP. • Solaris 2.4 or later.

Further information about IBM Connection Server SWs can be found at the URL http://www.internet.ibm.com.

3.5 Internet Servers

These are hardware platforms that contain preloaded software.

3.5.1 Internet POWERsolution for AIX - IBM Internet Connection Servers

This is a hardware/software combination of RISC/6000 hardware, an AIX operating system, and Internet Connection Server software. The software is pre-installed as an integral feature of the RS/6000 manufacturing process. The package can be connected by the customer to an Internet service provider. The POWERsolution IBM Internet Connection package consists of:

- RS/6000 model of customer's choice (except POWERparallel Systems and RISC system/6000 Model 40P - machine type 7020-all)
- AIX Version 4.1.4 or later
- Choice of IBM Internet Connection Server for AIX or IBM Internet Connection Secure Server for AIX software
- Sample home page library

3.5.2 Internet POWERsolution for AIX - Netscape Servers

This is a hardware/software combination of RISC/6000 hardware, an AIX operating system, and Netscape Server software. The software is pre-installed as an integral feature of the RS/6000 manufacturing process. The package can be connected by the customer to an Internet service provider. The POWERsolution Netscape package includes the following:

- RS/6000 7248 or 7024-E20
- AIX Version 4.1.4 or later
- Netscape Navigator Version 1.1 browser (comes with server)
- A choice of Netscape Communications Server Version 1.1 or Netscape Commerce Server Version 1.1 software
- Sample home page library

3.6 Lotus InterNotes

The InterNotes family of software products provide Web information and application integration between Lotus Notes and the Internet. This enables Lotus Notes users to publish Notes applications to the Internet and access the Internet directly from within Notes.

3.6.1 Lotus InterNotes Web Publisher

Creating, managing and updating enterprise Web servers is one of the biggest challenges that organizations face as they attempt to leverage the global reach of the Internet today. In most cases, Web sites are created and managed by a central group that gathers content from various contributors, manually converts that information into HTML, and creates the appropriate links. A very labor-intensive process to say the least.

The InterNotes Web Publisher specifically addresses the Web server challenge by leveraging the power of Notes′ distributed document creation and management system so that anyone throughout the organization that has access to Notes applications can automatically contribute to the company′s Web site(s).

The InterNotes Web Publisher is a Notes server program that automatically converts Notes documents and databases into HyperText Markup Language (HTML), the format used by standard Web browsers such as NCSA Mosaic and Netscape Navigator. Simply put, the InterNotes Web Publisher provides a simple, automated process for creating and managing Web sites.

Notes documents written by different people at different locations can quickly and easily be published to your Web site, obsoleting the need to manually re-create documents in HTML.

The InterNotes Web Publisher is a Notes Server application that runs in conjunction with a standard Web HyperText Transfer Protocol (HTTP) server. The InterNotes Web Publisher automatically converts Notes documents and views into a series of HTML documents that are accessible from a Web browser. By converting Notes views and doclinks, the InterNotes Web Publisher completely automates the process of creating and maintaining a navigable structure for a Web site.

When you publish a Notes database, the InterNotes Web Publisher does the following:

- Publishes the About Database document in the Notes database and makes it the home page for the database
- Lists the database views as hypertext links on the home page
- Converts each Notes document into an HTML file
- Converts Notes doclinks into hypertext links
- Converts Notes tables into HTML tables
- Converts bitmaps in Note documents into inline .GIF files
- Preserves attachments to Notes documents so users can download them from the Web site with a Web browser

3.6.1.1 Lotus InterNotes Web Publisher Administration

The primary interface to the InterNotes Web Publisher is the Configuration database, which resides on the Notes desktop. In this database, the administrator/Webmaster specifies:

- What databases to publish.
- Publishing interval (for example, every 2 minutes/hours/days, depending on the desired update cycle) for each database to be translated.
- Translation behavior: Do you want to publish all documents at each publishing interval, only publish those documents that have been added or modified, or remove the HTML from the Web site?

3.6.1.2 System Requirements and Configuration Options

The InterNotes Web Publisher runs on Windows NT and resides on a Notes server. It converts Notes databases to HTML files and places the resulting HTML files in a directory that should also be accessible to the HTTP server. The HTTP server can then make the files available to Web browsers, such as NCSA Mosaic and Netscape.

Basic requirements include:

- 486 or higher with 18 MB of RAM (32+MB recommended)
- 300 MB of disk space
- Microsoft Windows NT Advanced Server Version 3.1 or 3.5
- The Lotus Notes Server edition for Windows NT, Release 3.3 or higher
- A Web (HTTP) server
- An Internet connection

We recommend co-locating the Notes Server, InterNotes Web Publisher and HTTP server on the same machine for maximum performance.

Another possible configuration is to install the Internotes Web Publisher and the Notes server on one machine and the HTTP server on a separate machine, which is, in turn, connected to the Internet. The machine on which you install the HTTP server does not have to have Windows NT installed. However, the Web Publisher machine must be able to access the output directory to which the HTTP server software points.

Note: If your production Notes servers on your internal network are not on the Windows NT platform, simply replicate the databases you wish to publish from your production Notes server to the Notes server on NT for publishing. In addition, we strongly urge you not to connect your production Notes servers on your internal network directly to the Internet. Instead, use replication as a means of transferring information from your network to your Notes/InterNotes/Web machine, which will have the live Internet connection.

3.6.1.3 Lotus InterNotes Web Publisher 4.0

The InterNotes Web Publisher 4.0 enables businesses to create, manage, and administer their internal intranet and public Web sites using Lotus Notes Release 4. Businesses can use the proven application development facilities in Notes to easily build and host mission-critical applications on the Web. New InterNotes Web Publisher 4.0 features include leveraging Notes R4 for better Web content design and management, support for client and server-based imagemap creation, drag and drop building of Web views, improved search performance to Web clients and platform support for AIX, Sun Solaris and Windows 95 in addition to OS/2 and Windows NT.

InterNotes Web Publisher Release 4.0 automatically publishes Notes documents, views and forms to the Web, translating them into HTML. Businesses can take advantage of Notes′ collaborative authoring environment and workflow capabilities to automate the process of creating, approving, and consolidating Web content from multiple departments, ensuring a constant flow of up-to-date information to the Web site. In addition, Web content managed in the Notes document database is easy for Web browsers to navigate via Notes Views and is searchable using Notes′ full-text search engine.

Using InterNotes Web Publisher, any Web browser can participate in any Notes application (for example, lead generation, order taking, and customer service) by entering information into forms. Once the Web browser submits a form, InterNotes Web Publisher captures the information in a Notes database, enabling it to easily be incorporated into business process applications and core information systems.

For example, using InterNotes Web Publisher, businesses can easily create applications that enable Web browsers to request additional product information or a call from a salesperson. The Web browser simply fills out a form and submits it. The information in the form is captured in a Notes database where it can then be automatically routed for fulfillment purposes or added to existing information systems for later use.

3.6.1.4 System Requirements

InterNotes Web Publisher 4.0 requires a Notes Release 4.x server and any HTTP server with a TCP/IP connection. Platform support includes AIX, Sun Solaris, Windows 95, OS/2 and Windows NT. In Addition, InterNotes Web Publisher 2.1 (for use with Notes 3.x servers) is available on OS/2 and Windows NT platforms.

Further information about Lotus InterNotes Web Publisher such as pricing, versions availability, and download evaluation copy can be found at the URL `http://www.internotes.lotus.com`.

3.6.2 Lotus InterNotes News

Lotus InterNotes News 2.0 gives Notes users managed access to the newsgroup discussions that affect their business or industry. InterNotes News is a Notes server application that exchanges Usenet news articles between Notes and news servers; it uses the popular Internet standard Network News Transfer Protocol (NNTP), giving Notes users a secure and easy way to access and participate in Usenet newsgroups from the familiar Notes environment. By reading news articles contained in Notes discussion databases, users can leverage key Notes functionality, including hierarchical views of discussion threads, full-text search, and multiple indexed views of news articles.

3.6.2.1 Key Features

InterNotes News 2.0 offers users:

- An updated Newsgroup form. Buttons, such as Subscribe and Unsubscribe, have been replaced by Action buttons.
- Access to Usenet newsgroups without a personal Internet connection.
- Use of Notes agents, full-text search and mail forwarding to manage Usenet newsgroup articles.
- The ability to participate in newsgroups by writing and *posting* a response from Notes or by replying directly to the author using Notes mail (with an SMTP gateway).

InterNotes News 2.0 offers administrators:

- A choice between types of news feeds. Administrators can have news *pushed* to the InterNotes News Gateway or, for a more secure and controlled feed, they can *pull* news from a news server.

- Support for Notes R4 servers and clients. An upgraded database template for news messages and server process, supporting the 3-pane user interface in Notes R4.
- Options for configuring the cross post and spool interval. This new feature allows you to customize how often incoming articles are distributed among News databases and outgoing articles are sent to the INSPOOL.BOX.
- An updated News Database form. This allows administrators to turn off the creation of response hierarchies in news databases.
- Better performance. There is now support for running multiple InterNotes News processes.
- New console commands. To start a push transfer, issue the TELL INNEWS LISTEN command.
- The ability to make Internet newsgroups accessible to the organization without putting TCP/IP on every desktop.
- Controlled access and posting to newsgroups your organization deems appropriate.
- A centralized Notes configuration database that simplifies setup and administration of the News service. It allows administrators to subscribe to individual Usenet newsgroups, create customized Notes News databases and control News replication.
- Replication for easy distribution of news databases throughout the organization.

3.6.2.2 Platforms

InterNotes News supports R4 Notes servers running either Windows NT or IBM OS/2. Further information about Lotus InterNotes News can be found at the URL `http://www.lotus.com/webnews/`

3.7 Other Lotus Software Solutions to the Internet

- Lotus Domino Web Server
- Lotus Word Pro

3.7.1 Lotus Domino WebServer

Domino is a new server technology that transforms Lotus Notes into an Internet applications server allowing any web client to participate in Notes applications securely. Bridging the open networking environment of Internet standards and protocols with the powerful application development facilities of Notes, Domino provides businesses and organizations with the ability to rapidly develop a broad range of business applications for the Internet and intranet.

The majority of intranet/Internet sites today offer access to static information. Using Web technology as an information broadcast medium is merely the tip of the iceberg. Domino provides a rich set of facilities for building and hosting content-rich interactive Web sites. With Domino, businesses and organizations will realize the highest value from their Web investments as they use it to conduct business internally and externally.

Domino provides access to dynamic data and applications based on who you are to any Web client.

Domino provides all of the tools necessary to create and maintain content-rich interactive Web sites (the next wave of sites) through the power of Lotus Notes in conjunction with the open standards of the Web.

With Domino, you can create applications that leverage files stored in the file system of other Web servers or easily transmit and receive data from legacy systems.

Domino natively supports HTTP to render Notes data *on the fly* in HTML format, as well as to serve HTML documents from the file system. Using Domino, any Web client can now access and interact with Notes data and applications. For example, Web clients may create, edit and delete documents. Web clients benefit from the rich, collaborative applications, such as Sales Force Automation and Customer Service, developed and hosted in Lotus Notes.

In Addition, Domino takes advantage of Notes Access Control. Web site designers can deliver fine-tuned access control to Web sites and Web documents. Web users may be listed in the Notes Name and Address Book (Notes′ Directory Services), and when accessing a secured site, they are prompted for a valid name and password. The Web user′s access to functionality and information, down to the field level, is governed by predefined roles in the Notes Access Control List (ACL). In addition, Domino supports SSL, allowing server authentication and encryption of data at the session level.

Notes, combined with the Domino technology, provides the basic requirements for a Web site including a page management system, full-text search engine and threaded discussions. Coupled with Notes robust, rapid application development environment, it will enable customers to develop the next wave of Web sites hosting mission-critical business application.

3.7.1.1 Availability and Requirements

Domino beta is available for download from the World Wide Web at the URL `http://domino.lotus.com`.

Domino requires a Notes Release 4.x server.

Lotus Notes provides an ideal communications infrastructure by combining enterprise-ready, client/server messaging and the global access and distribution of the World Wide Web, together with a platform for rapidly developing and deploying strategic groupware applications. Notes enables individuals and organizations to communicate with colleagues, collaborate in teams, and coordinate business processes within and beyond their organizational boundaries to achieve improved business results. Lotus Notes supports all major operating systems: IBM OS/2 Warp, Apple Mac OS, UNIX platforms including IBM AIX, Sun Solaris, HP-UX, and SCO OpenServer, and Microsoft Windows and Windows NT. Notes is also available as a NetWare-loadable module for the Novell environment.

3.7.1.2 Domino Benefits

The following are the benefits associated with Domino:

- Reduces the complexity of creating and maintaining a content-rich Web site.
- Streamlines and automates the creation of content from multiple contributors/departments.

- Easy, graphical page management database reduces tedious links and creates a more navigable site.
- Eliminates the need to train content creators in HTML; anyone with word-processing capabilities can author Web site content.
- Gives Web application developers a rich environment for creating secure mission-critical interactive applications.
- Graphical forms designer.
- Broad range of application development facilities to serve the power user to the power programmer.
- Point and click creation of agents and formulas to advanced scripting capabilities.
- Integration with RDBMS and MQSeries.
- Integrated messaging system providing back-end infrastructure for business process applications workflow.
- Directory services for managing Web clients access to data and applications.
- Roles-based access control down to the field level.
- Domino provides all the facilities required to build a Web site:
 - Page management database.
 - Full-text search engine with automatic indexing of content.
 - Threaded discussion template.
 - Rapid application development of forms-based applications.
 - Registration template and directory services for secure Web client access.
 - Domino makes it possible to synchronously manage mirror sites and distributed intranets.
 - Secure and automated bi-directional synchronization of servers (replication) makes it easy to create mirror sites, distributed intranets, and update content and receive information from Web sites hosted by ISPs.

3.7.1.3 Domino Features

Domino makes it possible to use Notes′ rich application development environment to develop, manage and host Web applications.

Domino provides interactive Web client access to dynamic data and applications on a Notes server.

This means that Web clients may:

- Securely access a Notes server.
- Access dynamic data and application based on time, database queries and/or user identity.
- Create, edit and delete documents in a Notes database.
- Search a Notes database.
- View content in a Notes database with powerful Notes navigational capabilities such as the ability to expand and collapse views.

- Domino extends Notes Access Control to include Web clients.
- Updated template of Notes Name and Address Book form includes a new encrypted field to provide a Web client password.
- Web client authentication via Basic Web Authentication (name and password).
- Web user may be added to ACL lists, groups and rights and assigned a *role*.
- Database to field-level access control for Web clients.
- SSL support for server authentication and encryption of data in secured sessions.
- Domino serves HTML files stored in the file system.
- Domino runs CGI scripts activated by Web clients.

3.7.1.4 Internet/Intranet Applications

Domino provides businesses and organizations with the ability to rapidly develop a broad range of business applications for the Internet and intranet.

The following are some examples of applications:

- Customer service
- Sales automation: lead generation and tracking
- HR Benefits Program Information and Signup
- Threaded discussions for internal teams or for communities of customers

3.7.2 Lotus Word Pro

Lotus Word Pro is the first word processor to have direct Internet access and HTML editing built in. Direct Internet access means that you can use Word Pro to open a document from an FTP or Web server without having to first save it locally with a browser. Word Pro provides the tools that you need to create and edit HTML files and save them directly to the Internet without having to type cryptic tags and codes. Word Pro offers WYSIWYG editing, which means that what you see on the screen while you are creating your document is what people will see when they read your page with a Web browser such as Netscape.

Lotus Word Pro automates the common practice of editing and reviewing documents, enabling users to spend less time managing a team or tracking edits. Word Pro provides a step-by-step guide for setting a document up for review through Lotus′ TeamReview. Authors can easily assign access and editing rights for each individual reviewing a document, maximizing control over the editing process. Once multiple edits are made to a document, Lotus′ TeamConsolidate automates the process of consolidating these edits into one final draft. Word Pro is the only word processor that enables users to compare all edits on one screen instead of opening or printing multiple documents to view the edits. Users can make decisions interactively about which edits to accept or reject, thus shortening the editing time of collaborative documents.

In addition, Lotus Word Pro is the first word processor to provide document versioning technology, which enables users to store multiple versions of a document in a single file. Word Pro stores only the changes made between versions, making it an extremely efficient means for storage. Through versioning, users can track a document′s history and access previous versions. Versioning

not only maintains the integrity of each individual′s contributions to a document, but makes it easier for the document author to manage a document through its lifecycle.

3.7.2.1 Lotus Word Pro Redefines Word Processing

In rebuilding Word Pro from the ground up, Lotus approached even basic word processing functions from a new perspective. Among Word Pro′s innovative new concepts are SmartMasters, first seen in Freelance Graphics. The next generation of style sheets, SmartMasters contain *click here* blocks to guide users through the placement of information in a document, providing a powerful way to create professional looking documents. Unlike traditional templates, SmartMaster′s can also contain Divider Tabs. Similar to worksheet tabs in Lotus 1-2-3, Word Pro′s Divider Tabs provide an easy way to organize and navigate through long documents. Divider Tabs can correspond to parts of a document, such as the table of contents, chapter one, on chapter two, and can be contained in the document, linked to an external file, or linked to an OLE embedded object. Using Divider Tabs, users can store an entire work project in one file and share parts of a file with others on a team. Users can also drag and drop divider tabs to quickly rearrange information.

Further setting Word Pro apart from traditional word processors is its next generation spell check. In contrast to traditional spell checkers, Word Pro highlights all misspelled words at once and enables users to interactively spell check a document, significantly reducing editing time. Word Pro also allows users to mark text as a particular language, and quickly switch between English and any number of foreign language dictionaries.

Lotus Word Pro also features Lotus interface concepts, including the Task Sensitive Interface (TSI) and the Lotus InfoBox concepts, which make it easier for users to format and edit documents simultaneously.

3.7.2.2 Integration and Lotus Word Pro

The key to Lotus Word Pro is its ability to coexist with other word processing types including Word, WordPerfect, and DCA/RFT. Lotus Word Pro allows users to import a document from Word and WordPerfect, edit that document in Word Pro, and save it out in Word or Word Perfect without losing any formatting or data. Word Pro also supports both the SGML and HTML format, enabling users to easily create documents to be stored on the Internet without having to learn another package.

Lotus Word Pro is also tightly integrated with the Lotus family of products. Word Pro features LotusScript 3.0, Lotus′ cross-product object-oriented BASIC scripting language, and full OLE 2.0 support on Windows. Lotus Word Pro and Lotus SmartSuite share common code for features including spell check, SmartIcons, and routing. In addition, Lotus Word Pro features unique integration with Lotus Notes through technologies including Notes/FX, which facilitates the sharing of data between Lotus Notes and Word Pro.

3.7.2.3 Opening a File from the Internet

Opening a file from the Internet is as easy as opening it from your local hard drive. Once the file is open it looks just like it looked in your browser, without all of the confusing markup tags. Graphics on the page are displayed, as are tables and horizontal rules (lines). Even the background color is preserved.

3.7.2.4 Creating Your HTML Document

World Wide Web documents must be in HTML format so that browsers can display them and link them together. That's what enables you to click on a picture or a sentence and be taken elsewhere on the Internet. Traditionally, creating HTML files for the Internet required you to use an ASCII editor and type cryptic codes (known as HTML tags) around your words and sentences.

Word Pro includes a Smart Master, or template, that includes all of the character and paragraph tags that you will need to create your HTML file. Here's the list of tags that are literally at your fingertips:

- Address
- Anchor
- BlockQuote
- Citation
- Code
- Definition
- Emphasis
- Keyboard
- Preformat
- Sample
- Strong
- Typewriter
- Variable
- Default Text
- Definition Descriptions (1 through 5)
- Definition Terms (1 through 5)
- Example
- Heading (1 through 6)
- Horizontal Rule
- Ordered List (1 through 5)
- Unordered List (1 through 5)

Formatting your text is easy. All you do is choose the desired markup tag name from a list. For example, to create an ordered list, you would simply type your text, highlight it, then choose **Ordered List 1 (OL)** from the list of styles. Word Pro automatically numbers each item on the screen and puts in the and tags behind the scenes.

3.7.2.5 Converting Your Existing Files to HTML Documents

Do you have a collection of existing documents that you would like to publish on the Internet? Even if these documents are in other formats such as Frame Maker (MIF), Word 6, Word Perfect, etc., Word Pro can import them and convert them to HTML. Even your tables and graphics will be preserved. If your documents are structured with styles, you can map each style to an HTML tag so that all of your headings are automatically tagged as *Heading 1 (H1)*.

3.7.2.6 Creating Links (URLs) to Other Internet Documents

One of the best aspects of *surfing the net* is that after reading something, you can easily jump to a related topic by clicking on a word, sentence, or graphic. The hypertext reference code that allows this to happen is called a uniform resource locator (URL). Creating links in Word Pro is a simple process. The URL is typed into a Comment Note next to the text or graphic that will provide the link. A Comment Note in Word Pro is like an electronic post-it note or sticky note which can be hidden or displayed. By hiding the comment notes, the URL is still present, but it is hidden from your view so you′re seeing the document exactly as the browser will show it; also, the behind-the-scenes codes are not in your way.

3.7.2.7 Importing Graphics

Word Pro can import very many graphic formats, which are shown on the screen while you are editing your document. Graphic images can easily be moved around or resized by dragging them with the mouse.

Most Web browsers can only display graphics in JPEG and GIF format; thus, no matter what format your graphics were in when you imported them, Word Pro will automatically convert them to JPEG when you save your HTML file. The advantage is that you don′t have to convert each graphic by hand because Word Pro does it for you.

3.7.2.8 Tables

Word Pro supports HTML/2 format, plus several Netscape table extensions. Word Pro tables can contain connected cells and tables, graphics cells, and text within cells.

3.7.2.9 Saving to the Internet

After you′ve created your Web page or converted an existing document to HTML, you′re going to want to share it with the world. Word Pro can directly save files to FTP Host servers on the Internet (provided that you have the adequate rights to the server).

3.7.2.10 What Word Pro Needs to Exploit the Internet

If your computer is already set up to browse the World Wide Web, then it is ready for Word Pro. You must have an active TCP/IP connection to use the FTP and HTTP (World Wide Web) clients built in to Word Pro. Word Pro works with any WinSock-compliant TCP/IP protocol stack that connects via SLIP or PPP dialers or through corporate proxies or firewalls. There is nothing to configure in Word Pro unless you are accessing the Internet through a corporate proxy or firewall.

For more information on HTML editing, go to the URL `http://www.ncsa.uiuc.edu/demoweb/html-primer.html`.

3.7.2.11 Availability and Software Requirements

Lotus Word Pro is available to the Windows 3.1, Windows 95 and OS/2 platforms.

System requirements for Lotus Word Pro, which is currently in beta testing, are estimated at a minimum 386 IBM or compatible PC with 33 MB of hard disk space and 6 MB of RAM.

Further information about Lotus Word Pro can be found at the URL `http://www.lotus.com/wordpro/`.

Chapter 4. Web Development

When you're going to develop home pages, one of the first things you have to consider are the platforms you have to use, the language you use, the interfaces, and the databases, and you have to integrate them in a heterogeneous environment. If you choose a database system like DB/2 and make an application outside the web (Internet or intranet), you have to be sure that all the people that have to use it have the DB/2 client installed on their system. This requirement magnifies the problem if the users have different operating system environments (such as (AIX, Solaris, OS/2, Windows 3.x, DOS, Windows 95, Windows NT, etc. You'll have to seek a client for all the machines, a good number of licences and so on. You'll also have to work on migrating your job to each platform.

If you use the DB2 WWW gateway you'll have to buy explorers for each platform (don't bother if the browsers are from different companies) and make your DB2 WWW macros on your Web server. So now you're ready to go. You only had to write it once and you didn't have to worry about the platform, the compilation, etc. So you must develop home pages in order to improve your network and application flexibility.

The first thing you need to know about Web development is how to make pages. Once you make your interface with your home pages, develop the interfaces with the final objectives (databases, mail, or just plain text files). You can do this with the help of 2 tools: CGIs and Java. Finally, the initial work is done (feedback is always very important, a system is something that is never finished).

4.1 Hypertext Markup Language (HTML)

The HyperText Markup Language (HTML) is the language used to write hypermedia documents for the World Wide Web (WWW). HTML is a subset of the Standard Generalized Markup Language (SGML); SGML is an international standard for document markup conforming to ISO 8879.

The latest defined version of HTML is HTML3.0.

HTML is similar to a computer programming language; there are commands called tags and syntax rules to be observed when writing in HTML.

HTML documents can be written using any word processor or text editor. However, the way they look when seen with a Web browser is quite different from what the writer sees when editing them; it is not the what you see is what you get (WYSIWYG) approach. Some WYSIWYG HTML editors are currently available and will be covered later in the chapter.

The HTML language provides support for the following features:

- Hypertext links to resources (documents, multimedia or data files)
- Menus and forms
- In-line graphics
- Text formatting

© Copyright IBM Corp. 1996

4.1.1 HTML2.0 Document Structure

HTML documents are composed of two main parts: a head and a body. Every HTML document should start with a head. The head is the top part of the document; it generally includes the document's title. Different browsers use different ways to display the document's title. NCSA Mosaic, for instance, displays it in a field named Document Title right under the menu bar, while WebExplorer displays it in the title bar. The title is also the way by which documents are referenced when saved in the Hotlist or Quicklist of the browsers. It should therefore be short enough to fit into one line of the Hotlist window but still be descriptive descriptive. An optimized title length is around 64 characters. Besides the title, document heads can contain information about the document type. Index documents, for instance, are identified in the head as such documents. The head of a document cannot contain anchors, any kind of highlighting or paragraphs. The head of the document is enclosed between a <HEAD> and a </HEAD> tag.

The second main part of an HTML document is the body. The body is the core part of the document; it contains all the information that is part of the document and controls the way this is presented to browser users. The body can contain images, links to other resources, lists, menus, entry fields, or plain text. The body of the document is enclosed between a <BODY> and a </BODY> tag.

4.1.2 HTML2.0 Syntax

The HTML language uses markup tags to identify the elements of the documents. All tags begin with a left angle bracket (<) and end with a right angle bracket (>). Except for a few, all tags are containers. This means that there's always an opening tag and a closing tag. For example, an unordered list is opened by and closed by . The following table contains the main HTML elements:

Table 18 (Page 1 of 2). HTML Main Elements

Name	Opening tag	Closing tag	Description
Anchor	< A >	< / A >	HyperLink to a resource
Address	<ADDRESS>	</ADDRESS>	Format an address
Bold	< B >	< / B >	Display text in bold
Base	<BASE>	no closing tag	Record URL of document
Body	<BODY>	</BODY>	Contain the document's body
Blockquote	<BLOCKQUOTE>	</BLOCKQUOTE>	Include text in quotes
Line Break	< B R >	no closing tag	Break current line
Citation	<CITE>	</CITE>	Specify a citation
Code	<CODE>	</CODE>	Enclose an example of code
Definition list description	< D D >	no closing tag	Description of definition list item
Directory list	<DIR>	</DIR>	Enclose a directory list
Definition list	< D L >	< / D L >	Enclose a list of terms and definitions
Definition list item	< D T >	no closing tag	Item of definition list
Emphasis	< E M >	< / E M >	Emphasize enclosed text
Form	<FORM>	</FORM>	Define form of enclosed text

Table 18 (Page 2 of 2). HTML Main Elements

Name	Opening tag	Closing tag	Description
Level 1 heading	<H1>	</H1>	Enclose level 1 heading
Level 2 heading	<H2>	</H2>	Enclose level 2 heading
Level 3 heading	<H3>	</H3>	Enclose level 3 heading
Level 4 heading	<H4>	</H4>	Enclose level 4 heading
Level 5 heading	<H5>	</H5>	Enclose level 5 heading
Level 6 heading	<H6>	</H6>	Enclose level 6 heading
Head	<HEAD>	</HEAD>	Define the head of the document
Horizontal rule	<HR>	no closing tag	Insert horizontal line
HTML	<HTML>	</HTML>	Define HTML document
Italics	<I>	</I>	Italicize enclosed text
Image		no closing tag	Embed an image
Input	<INPUT>	</INPUT>	Display entry field
Index	<ISINDEX>	no closing tag	Define searchable URL
Keyboard	<KBD>	</KBD>	Indicate user typed text
List item		no closing tag	Item of directory list, menu list, ordered list, unordered list
Link	<LINK>	no closing tag	Describe relationship between documents
Menu	<MENU>	</MENU>	Enclose a menu list
Ordered list			Enclose an ordered list
Option	<OPTION>	no closing tag	Indicate one choice in a select menu
Paragraph	<P>	</P>	Define a paragraph
Preformatted text	<PRE>	</PRE>	Enclose preformatted text
Sample	<SAMP>	</SAMP>	Indicate sample text
Select	<SELECT>	</SELECT>	Define a set of selectable options
Strong emphasis			Strongly emphasize text
Title	<TITLE>	</TITLE>	Define document's title
Typetype	<TT>	</TT>	Display enclosed text in monospaced font
Textarea	<TEXTAREA>	</TEXTAREA>	Enclose a text area
Underlined	<U>	</U>	Underline text
Unordered list			Enclose an unordered list
Variable	<VAR>	</VAR>	Indicate a variable

HTML tags are case insensitive; every command is interpreted by the browsers independent of the capitalization; the tag <SELECT>, for example, can either be written <Select>, <select>, or <sELecT> without making any difference.

The most commonly used HTML tags are the Headings, Lists, Anchors or Links, Images and Forms tags.

4.1.2.1 Headings

HTML supports up to six heading levels; their tag is <H*>, where * is a number from one to six. Headings change the font of the embedded text, put breaks before and after it and render the text. Figure 17 shows how the six HTML heading levels are rendered by a Web browser. The Web browser that is shown in the figures of the current chapter is WebExplorer, the OS/2 Web browser.

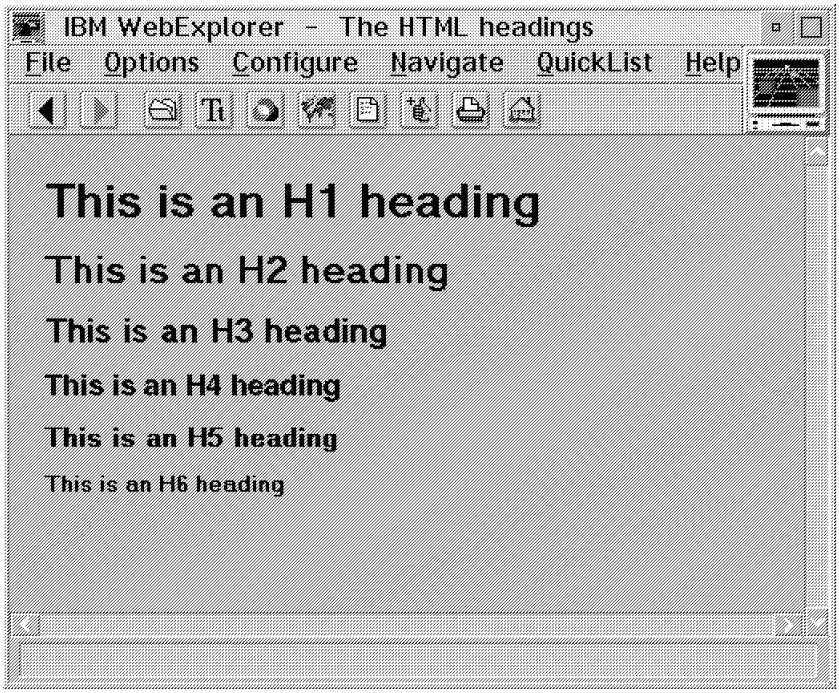


Figure 96. HTML Headings. Web browser rendering of the six HTML heading levels.

4.1.2.2 Lists

Lists are heavily used in the body of HTML documents. They are basically containers that include items; in this section we will show how to write lists in HTML and how these lists are displayed by browsers. There are five supported types of lists; they are:

- Definition List
- Directory List
- Menu List
- Ordered List
- Unordered List

Definition List: The following is an example of a definition list:

```
<DL>
<DT> First item <DD>First item's definition
<DT> Second item <DD>Second item's definition
<DT> Third item <DD>Third item's definition
</DL>
```

Figure 97 shows how the definition list is displayed by the Web browser.

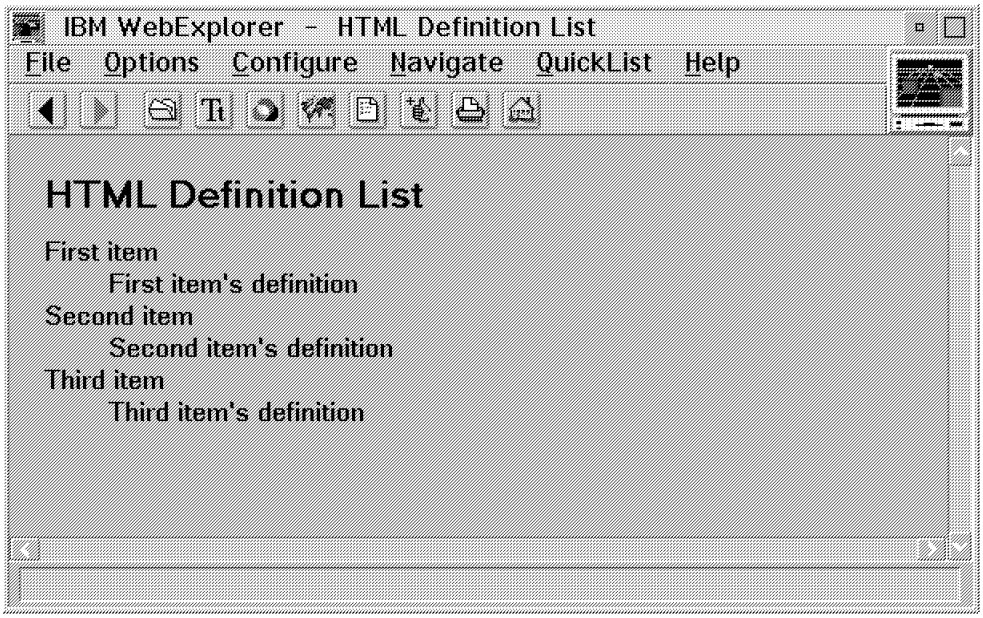


Figure 97. HTML Definition List. Web browser rendering of an HTML definition list.

Definition lists can have the COMPACT attribute. In this case they are rendered with a reduced width.

Directory List: The following is an example of a directory list:

```
<DIR>
<LI>A-L
<LI>M-R
<LI>S-Z
</DIR>
```

Figure 98 shows how the directory list is displayed by the Web browser.

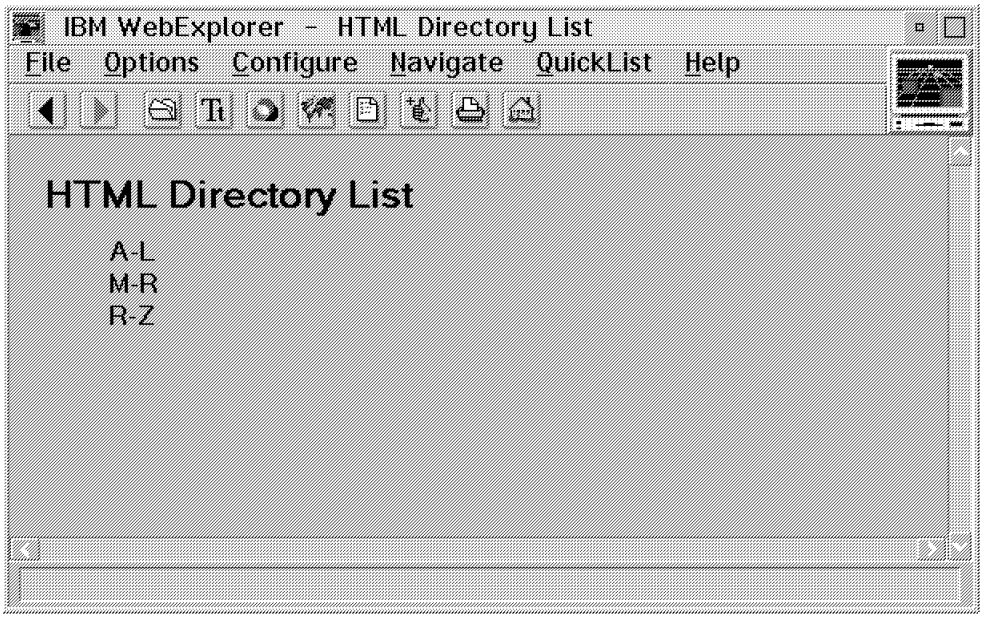


Figure 98. HTML Directory List. Web browser rendering of an HTML directory list.

Menu List: The following is an example of a menu list:

```
<MENU>
<LI>First menu item
<LI>Second menu item
<LI>Third menu item
</MENU>
```

Figure 99 shows how the menu list is displayed by the Web browser.

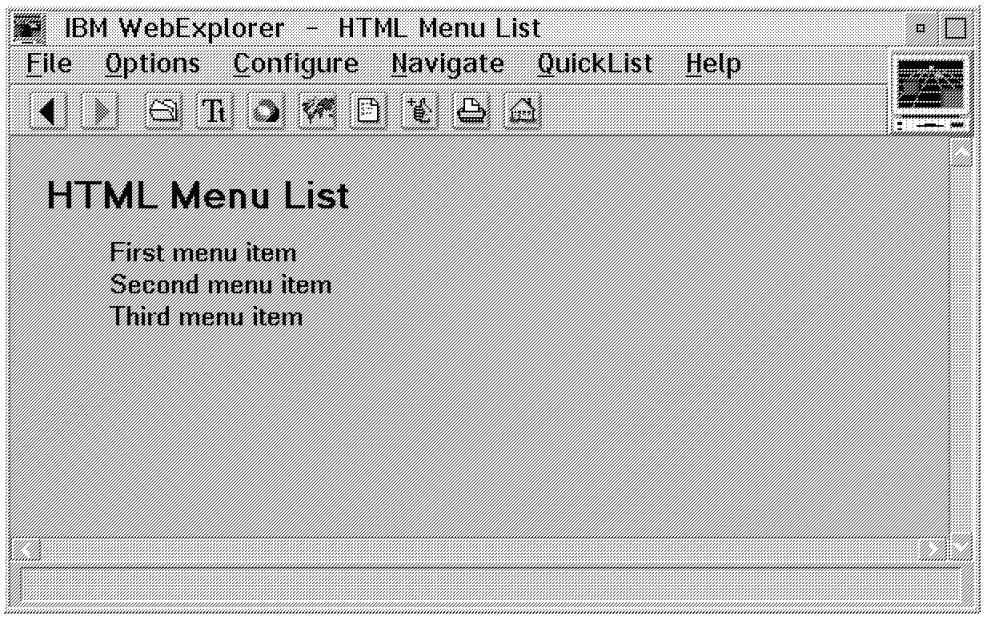


Figure 99. HTML Menu List. Web browser rendering of an HTML menu list.

Ordered List: The following is an example of an ordered list:

```
<OL>
<LI>First list item
<LI>Second list item
<LI>Third list item
</OL>
```

Figure 100 shows how the ordered list is displayed by the Web browser.

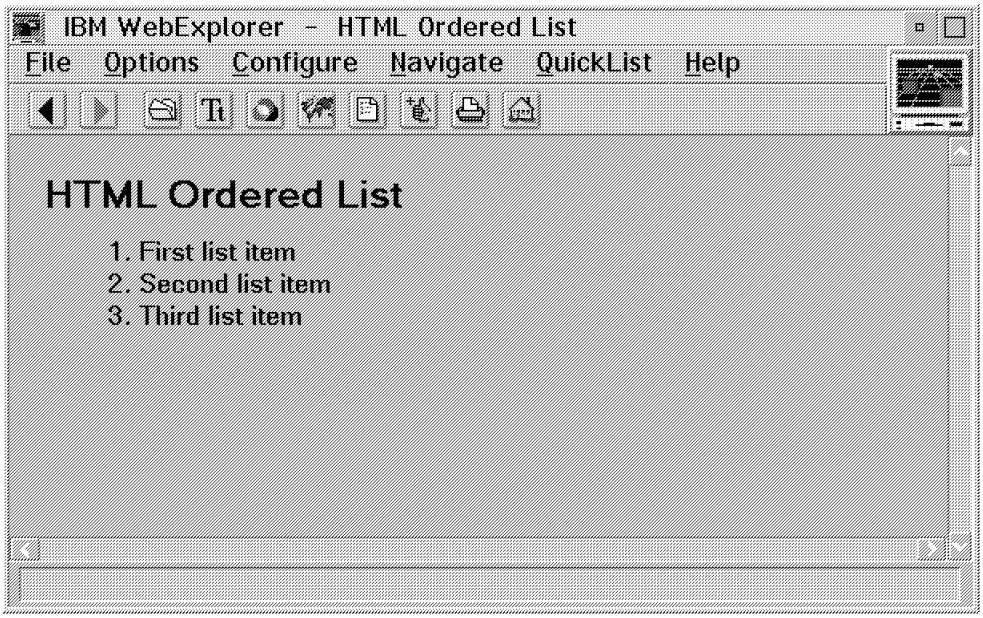


Figure 100. HTML Ordered List. Web browser rendering of an HTML ordered list.

Unordered List: The following is an example of an unordered list:

```
<UL>
<LI>First list item
<LI>Second list item
<LI>Third list item
</UL>
```

Figure 101 shows how the unordered list is displayed by the Web browser.

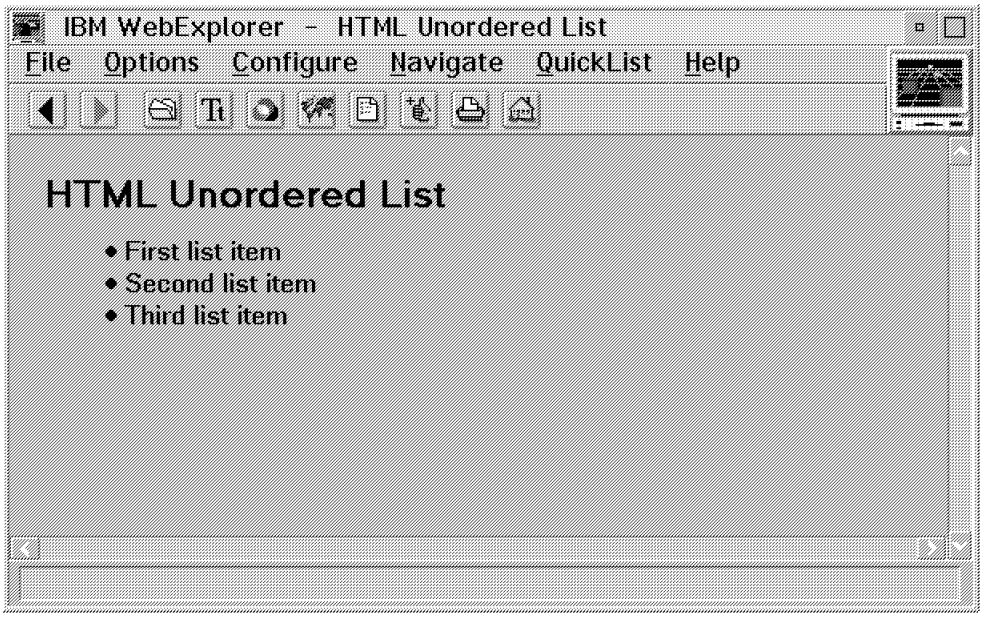


Figure 101. HTML Unordered List. Web browser rendering of an HTML unordered list.

4.1.2.3 Anchors

Anchor tags specify links to resources available on other systems or somewhere else on the local system. Links can be represented by text or images. In the first case, the text is in hypertext and the link is a hypertext link; in the second case, the link is an image link.

The link is activated by clicking on the hypertext or the image. This will cause the Web browser to retrieve the linked document and display it in place of the one currently displayed. Web browsers show hypertext links in a different color than normal text. When the mouse pointer is positioned over a hypertext link or an image link, its pointer's icon changes to indicate that clicking the mouse button will activate the link.

Anchors are identified by the <A> tag and their syntax is as follows:

```
<A HREF="URL">Hypertext</A>
```

URL is the Uniform Resource Locator of the pointed resource.

The URL (Uniform Resource Locator) points to a resource that can be on any machine on the Internet. The pointed resource is not necessarily another HTML file; it may be any other kind of file or it may not even be a file. It could be the result of a database query. The serving protocol specified by the URL is not necessarily HTTP. It can be any one of the following:

- HTTP
- Gopher
- WAIS
- FTP
- File
- News

The following example shows an HTML anchor that creates a hyperlink to a Home Page located on the www.austin.ibm.com server:

```
<A HREF="http://www.austin.ibm.com/Home.html">IBM Austin Home Page</A>
```

The text (IBM Austin Home Page) that is between the <A> and tags is what will be displayed as the hyperlink when this anchor is displayed by a browser. When the reader clicks on this text, the browser will load the Home Page referred to by the URL in the anchor.

Hyperlinks do not necessarily have to be other Web Pages; they can be, for example, Gopher or Telnet connections. The following example of an HTML anchor shows how to create a link to a Gopher server:

```
<A HREF="gopher://gopher-vm.almaden.ibm.com">Almaden Gopher Server</A>
```

This example shows how to create a link to a Telnet server:

```
<A HREF="telnet://telnet.w3.org>A telnettable browser<A>
```

Anchors can also be used to create hyperlinks to HTML files that are stored on the reader's local system. For example:

```
<A HREF="Catalog.html">Catalog</A>
```

In this case, the file Catalog.html is an HTML file that is on the readers local system.local file. The browser will resolve the URL to:

```
http://our.host.com/Catalog.html
```

where our.host.com is the reader's system host name.

Another way to use an anchor is to make it point to another place in the current document; this is shown in the following example:

```
<A HREF="#Info">Information</A>
```

The Information hyperlink, when selected, will branch to a location in the currently displayed HTML file that has the associated anchor point. This anchor point would be specified with the following HTML anchor:

```
<A NAME="Info">Information</A>
```

It is also possible to point to an anchor point in another document, as shown in the following example of an HTML anchor:

```
<A HREF="http://remote.host.com/Info.html#Info">Information </A>
```

The anchor point is specified in the document referenced by the anchor's URL in the same way as the anchor point is specified earlier. Specifically:

```
<A NAME="Info">Information</A>
```

4.1.2.4 Images

HTML documents can imbed images and control their position and the position of the text beside them. Import of images is tagged with and can have the following parameters:

- SRC=URL to define the link to the image file
- ALIGN=TOP, MIDDLE, or BOTTOM to define the position of the text next to the image
- ALT, alternative text to be displayed in a nongraphic environment
- ISMAP, to make the image a map

The following is an example that causes a GIF format image that resides on the reader's local system to be displayed:

```
<IMG SRC="image.gif">
```

If the image file is located somewhere else on the Internet, the syntax for the link would be:

```
<IMG SRC="http://remote.host.com/image.gif">
```

The ALIGN parameter determines the position of the text beside the image. It can assume three values; the following examples show their results:

ALIGN=TOP. The text is positioned at the top of the image. Here is an example of an HTML statement that imbeds an image in the document using this option:

```
<IMG ALIGN=TOP SRC=pmglobe.gif"> Globe image
```

Figure 102 shows how this option is displayed by the Web browser.

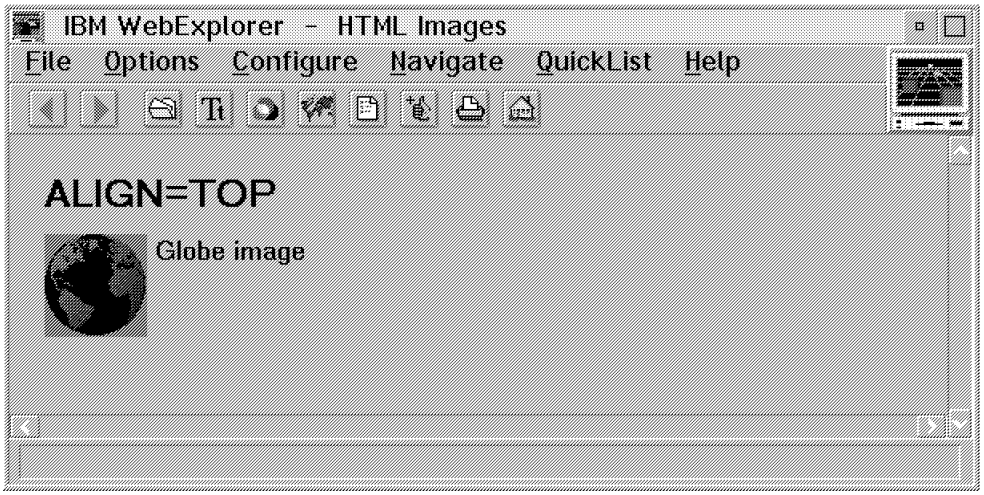


Figure 102. HTML Figures. Web browser rendering of the text beside a figure when ALIGN=TOP is chosen.

ALIGN=MIDDLE. The text is positioned at the middle of the image; the following is an example of the HTML statement:

```
<IMG ALIGN=MIDDLE SRC=pmglobe.gif"> Globe image
```

Figure 103 shows how this option is displayed by the Web browser.

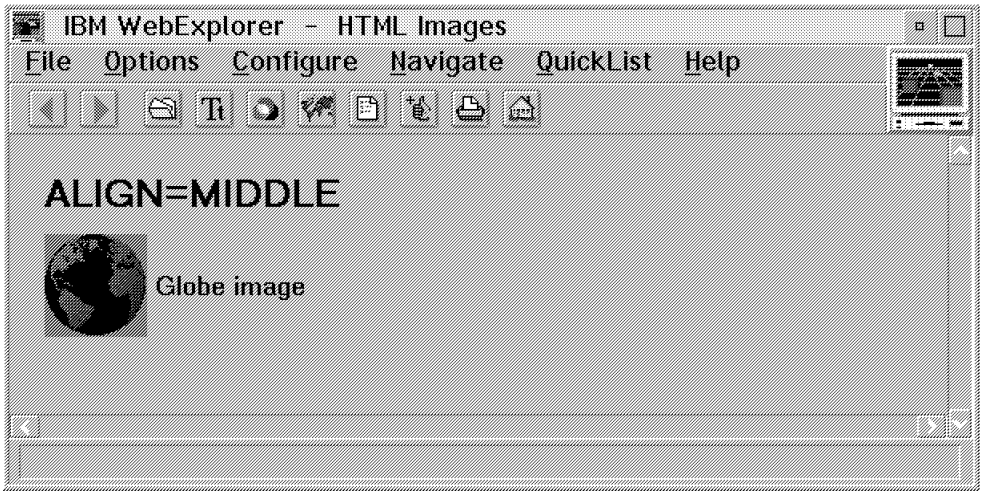


Figure 103. HTML Figures. Web browser rendering of the text beside a figure when ALIGN=MIDDLE is chosen.

ALIGN=BOTTOM. This is the default. The text is positioned at the bottom of the image; the HTML statement is:

```
<IMG ALIGN=BOTTOM SRC=pmglobe.gif"> Globe image
```

Figure 104 shows how this option is displayed by the Web browser.

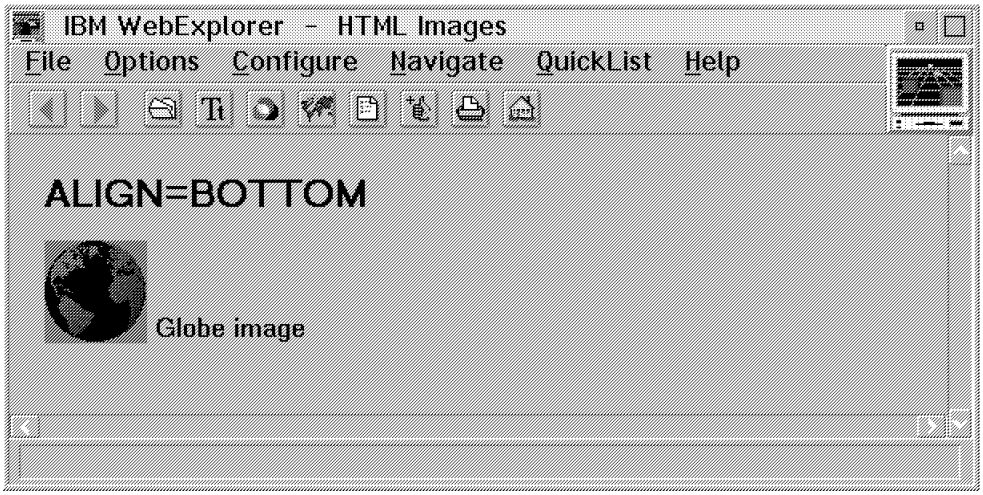


Figure 104. HTML Figures. Web browser rendering of the text beside a figure when ALIGN=BOTTOM is chosen.

Although there is not a tag that allows you to indent images on a Web Page, images can be shifted to the right using the preformatted text <PRE> tag followed by a number of blanks; the following is an example:

```
<PRE>
<IMG SRC="pmglobe.gif"> Globe image
</PRE>
```

Figure 105 shows how this is displayed in the Web browser.



Figure 105. HTML Figures. Web browser rendering of a figure shifted to the right using the <PRE> tag.

Note: Be aware that by using the <pre> tag, you have a type face shift for the text associated with the image. With the inclusion of the </pre> tag you will return to the regular type face.

An image can also be a link to another document. To make this happen, the tag is imbedded within an anchor, such as the following:

```
<A HREF=http://remote.host.com/Homepage.html><IMG SRC="pmglobe.img"></A>
```

In this example, the image itself is the hyperlink to the Home Page specified by the URL in the anchor. When this anchor is displayed by a Web browser and the reader moves the mouse pointer over the image, the mouse pointer icon changes the same way it does when it's positioned over a normal text hyperlink. Whenever an image is used as a hyperlink, it is important that the hyperlink also contain some text. This allows people using text-only displays to still hyperlink to the specified resource. For example:

```
<A HREF=http://remote.host.com/Homepage.html>
<IMG SRC="pmglobe.img">My Home Page</A>
```

4.1.2.5 Image Maps

Using the parameter ISMAP, the image is transformed into a map. An Image Map is a particular kind of image that, when displayed by a Web browser, is able to sense the position of the mouse pointer on itself. It is then possible to make different portions of the same map point to different resources.

The image that is displayed is the same kind of image that is used for normal images, with the addition of some more information that is needed to cause the browser to be able to sense the mouse pointer position. The procedure to be followed to set up a map on a WWW server depends on the server software installed on the machine. Here we describe the procedure for an NCSA HTTP server; for other servers, refer to the related documentation.

Image Maps are set up as follows:

1. The server that is to serve the Image Map must be configured to support Image Maps. This is done by:

 - Compiling the imagemap program located in the cgi-src directory with the command:

     ```
     make imagemap
     ```

 - The imagemap program uses a configuration file that is located in /usr/local/etc/httpd/conf/imagemap.conf. If you would like to change the location of this file, edit cgi-src/imagemap.c, change the setting of CONF_FILE, and recompile with the command:

     ```
     make imagemap
     ```

2. The image to be used as an Image Map must be created as a GIF format file. It can be created with drawing tools, screen capture utilities or any program that can generate a GIF format file. This procedure uses an example image named mapimage.gif.

 Even though any GIF image can be turned into a map, it makes more sense to use pictures that contain sharply separated elements so the users can easily tell which part of the image they are pointing at with the mouse.

3. An Image Map configuration file must be created that establishes the links between portions of the image and other resources. The easiest way to divide the image is to split it into portions of rectangular shape. This procedure uses an example Image Map configuration file called /mapdir/mapfile.map. This example Image Map configuration file, shown below, divides our example image into four separate rectangles, each linking to a different resource.

```
default /X11/mosaic/public/local.html

rect (12,10) (70,30) http://first.link.com/first.html
rect (80,40) (100,50) http://secondf.link.com/second.html
rect (120,70) (170,100) ftp://third.link.com/
rect (200,100) (250,150) http://fourth.link.com/fourth.html
```

The first statement in the file defines the default link. This is the one to be hyperlinked when the user clicks with the mouse on an area of the map that doesn't belong to any of the rectangles defined below. In this example, the default points to a local file named /X11/mosaic/public/local.html.

The remaining statements define the links between rectangular areas of the image and the hyperlink resources; In this example, the first set of coordinates establishes a hyperlink to the URL http://first.link.com/first.html. This document will be hyperlinked whenever the user clicks the mouse button within the area of the Image Map contained in the rectangle whose upper left corner has pixel coordinates 12,10 and whose lower right corner has pixel coordinates 70,30.

This example used a rectangular shape. The various supported shapes and their syntax are:

- rect (x·, y·) (x·, y·) URL

 Defines a rectangle's upper-left and lower-right corner coordinates in pixels.

- circ (x·, y·) r URL

 Defines a circle by it's center's coordinates and radius in pixels.

- poly (x·, y·) (x·, y·)(x_n, y_n) URL

 Defines a polygon by giving the coordinates of its vertices in pixels.

The coordinates of the image specified in this file can be found using any good graphic editor.

4. The server's imagemap.conf configuration file mentioned earlier must be modified to include an entry that establishes a name for the Image Map file previously created. For example,

   ```
   mymap  :  /mapdir/mapfile.map
   ```

 mymap - This is any name, that you desire, which will be used to reference the Image Map configuration file. This example uses the name mymap.

 /mapdir/mapfile.map - This is the full path file name of the Image Map configuration file.

5. The last step is to add an HTML anchor for the Image Map in your HTML document. For example:

   ```
   <A HREF="http://machine/cgi-bin/imagemap/mymap">
   <IMG SRC="mapimage.gif" ISMAP>
   </A>
   ```

 machine - This is the name of the server which is to serve the Image Map.

 mymap - This is the name that you called the Image Map's configuration file in the imagemap.conf file.

 mapimage.gif - This is the name of the GIF image.

 There is no limit to the number of Image Maps that a Web server can serve.

Information on the Image Map creation for other HTTP servers, as well as further information on the NSCA server, can be found at the following URL:

http://www.w3.org/hypertext/WWW/Daemon/User/CGI/HTImageDoc.html

Examples of Image Maps can be found on the Web at the following URLs:

http://wings.buffalo.edu/world/
http://www.nchcp.lcs.mit.edu/Info/structure.html
http://www1.cern.ch/Demo/Images/Dragons.html
http://www.hcc.hawaii.edu/hccinfo/hccmap/hccmap2.html

4.1.2.6 Forms

Forms are parts of an HTML document that allow the reader to input information that will be sent back to the server for processing. You can define many Forms in a single document. However, Forms cannot be nested. In other words, you cannot put a Form within a Form. Each Form can contain interactive elements, such as text input fields, push buttons, radio buttons, check boxes, and option menus. These elements are used to request information from the reader. When the reader enters the requested information, their information is sent back to the server and processed by a CGI script.

A Form is constructed by including a <FORM> tag and one or more Form definition tags in an HTML document. There are actually five tags that are used to define a Form: a <FORM> tag and four Form definition tags. These tags are:

<FORM> Define a form

<INPUT> Define an input field

<OPTION> Define selectable options

<SELECT> Define a list of selectable options

<TEXTAREA> Define a multiline input field

Each one of these tags can have attributes that define in more detail the characteristics of the Form. Let's look at each one of these tags in more detail.

- <FORM> Tag

 The <FORM> tag defines the overall characteristics of the Form and delimits the Form definition tags that define the contents and layout of the Form. The <FORM> tag can have the following attributes:

 ACTION Specifies the URL of the address of the server and CGI script that will process the reader's input to the Form.

 METHOD Selects the method that the server will use to pass the reader's input to the CGI script. Its values can be GET and POST; the first one puts the Form data into a CGI environment variable, and the second passes it to the CGI script as standard input (stdin).

 ENCTYPE Specifies the encoding for the Form input. This attribute only applies if METHOD is set to POST and is rarely used.

 The <FORM> tag always requires the closing tag </FORM>.

- <INPUT> Form Definition Tag

 The <INPUT> tag defines an input field on the Form. This tag can have several attributes which define the name of the field, its layout, the type of input, maximum input length, and range of acceptable input values. These attributes are:

ALIGN Used to specify the vertical alignment of the image when TYPE=image. The values that it can assume are TOP, MIDDLE and BOTTOM.

CHECKED A flag that which indicates the radio button or checkbox being defined by this INPUT tag is initially selected.

MAXLENGTH Indicates the length of the field in characters.

NAME Symbolic name of the variable to which the input field value is assigned.

SIZE Specifies the size of the field according to its type. The number assigned to it is the length in characters of the visible part of the field.

SRC URL or URN of the image. Used only if TYPE=image.

TYPE Defines the type of input field. Although HTML tags are supposed to be case insensitive, some browsers do not display the form correctly if the values of the TYPE parameter are capitalized.

checkbox Used for boolean or for multiple selectable choices.

hidden No visible input field, but its content is sent with the form.

image Define the image field to click on with the mouse to submit the Form.

password Input text not to be displayed when entered.

radio Used for mutually exclusive choices.

reset Defines a button that, when pressed, resets fields to their initial values.

submit Defines a button that, when pressed, submits the Form.

name Name of the submitted data.

text Defines a single-line entry field.

VALUE Value to be returned when a field is selected or an initial value is displayed in the field.

The <INPUT> tag has no closing tag.

- <OPTION> Form Definition Tag

 The <OPTION> tag is used in conjunction with the <SELECT> tag to define an option dialog. One or more <OPTION> tags are specified for each <SELECT> tag to define the options that the user has to choose from. The <OPTION> tag can have the following attributes:

 DISABLED The choice is not selectable.

 SELECTED Indicates the initially selected choice. If it is not specified, the first item of the list is initially selected.

 VALUE The value to be returned if the option specified by this tag is chosen.

 The <OPTION> tag has no closing tag.

- <SELECT> Form Definition Tag

 The <SELECT> tag is used in conjunction with the <OPTION> tag to define an option dialog. The <SELECT> tag defines the characteristics of

the option dialog and delimits the <OPTION> tags that are used to specify the available option choices. The option dialog will be displayed differently depending on the browser the reader is using. However, it is normally displayed as a pull-down list, a pop-up list or a scroll list. The <SELECT> tag can have the following attributes:

ERROR Used to indicate that the initial selection is in some way in error.

MULTIPLE Allows the reader to make multiple selections from the dialog. The default is that only one selection is allowed.

The <SELECT> tag always requires the closing tag </SELECT>.

- <TEXTAREA> Form Definition Tag

 The <TEXTAREA> tag is used to define a multiline input field. <TEXTAREA> has the following attributes:

 ROW Number of rows in the input field.

 COLS Number of columns in the input field.

 The <TEXTAREA> tag always requires the closing tag </TEXTAREA>. Figure 106 shows how a document containing a Form is displayed by a Web browser.

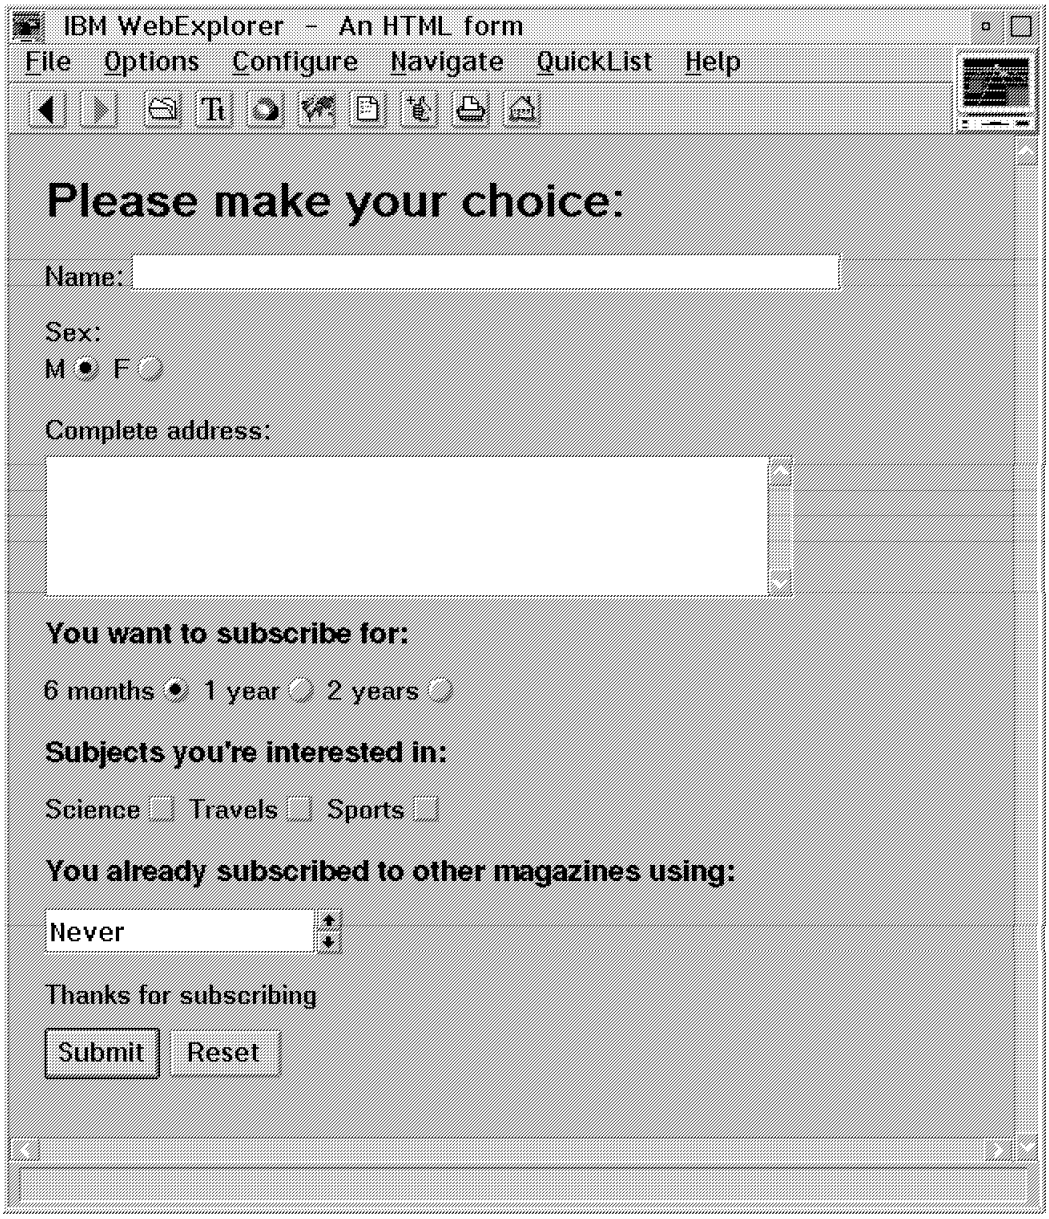


Figure 106. HTML Form

The HTML source for the document shown in Figure 106 is the following:

```
<HTML>
<HEAD>
<TITLE>
An HTML form
</TITLE>
</HEAD>
<BODY>
<H1>
Please make your choice:
</H1>
<FORM METHOD="GET" ACTION="http://WebServer/cgi-bin/mailit.pl">
<P>Name: <INPUT NAME="name" SIZE="36">
<P>Sex: <BR>
M     <INPUT NAME="sex" VALUE="m" TYPE=radio>
F     <INPUT NAME="sex" VALUE="f" TYPE=radio>
<P>Complete address:
<TEXTAREA NAME="address" COLS=36 ROWS=4>
</TEXTAREA>
<H4>You want to subscribe for: </H4>
6 months <INPUT TYPE="radio" NAME="sub" VALUE="1">
1 year <INPUT TYPE="radio" NAME="sub" VALUE="2">
2 years <INPUT TYPE="radio" NAME="sub" VALUE="3">
<H4>Subjects you're interested in: </H4>
Science <INPUT TYPE="checkbox" NAME="top" VALUE="5">
Travels <INPUT TYPE="checkbox" NAME="top" VALUE="6">
Sports <INPUT TYPE="checkbox" NAME="top" VALUE="7">
<H4>You already subscribed to other magazines using: </H4>
<SELECT NAME="alr">
<OPTION SELECTED>On line forms
<OPTION>Phone
<OPTION>Mail
<OPTION>Other
</SELECT>
<P>Thanks for subscribing
<P><INPUT TYPE=submit> <INPUT TYPE=reset>
</FORM>
</BODY>
</HTML>
```

The line from the form that reads:

```
<FORM METHOD="GET" ACTION="http://WebServer/cgi-bin/mailit.pl">
```

This specifies the URL of the CGI script that will process this Form. In this example, the PERL script mailit.pl in the cgi-bin directory on the Web server named WebServer will process the form.

For more information about HTML Forms, see the following URL:

```
http://www.yahoo.com/Computers/World_Wide_Web/Programming/Forms/
```

4.1.3 HTML3.0 or HTML+

As HTML was used to publish information on the Web, some limitations in its capabilities were found. HTML, for example, is not able to enclose mathematical formulas or tables of any kind in its documents. From a performance viewpoint, retrieving large documents from a server takes time, and HTML was not designed with the capability to split large documents over several servers.

To address these problems a new language emerged. This new language is called HTML+ or HTML3.0. It is an enhancement of HTML and was designed to address the problems found with HTML by adding new capabilities to the HTML language. As of the publish date of this redbook, the HTML+ specifications were still in draft form; the final documentation should be available shortly thereafter. The following information is then based on draft specifications and slight changes might be necessary in the future. Some HTML tags have been dropped and included as attributes of other tags. Backward compatibility with HTML documents is assured. However, simple programs are available to convert HTML documents into HTML+.

The major enhancements of HTML+ over HTML are:

- Major changes to <BODY> tag
- Split large documents across multiple servers
- Support for tables
- Support for mathematical formulas

The document's structure is basically the same as HTML. The two main parts of a document are the heading and the body. More control tags have been introduced in HTML+ to support its enhanced features; the following is a table listing these features:

Table 19 (Page 1 of 2). HTML+ New Elements

Name	Opening tag	Closing tag	Description
Abbreviation	<ABBREV>	</ABBREV>	Enclose abbreviations
Abstract	<ABSTRACT>	</ABSTRACT>	Enclose abstracts
Acronym	<ACRONYM>	</ACRONYM>	Enclose acronyms
Added	<ADDED>	</ADDED>	Enclose added text
Argument	<ARG>	</ARG>	Enclose arguments
Array	<ARRAY>	</ARRAY>	Define mathematical matrices
Box	<BOX>	</BOX>	Group mathematical items
Byline	<BYLINE>	</BYLINE>	Info on document authors
Caption	<CAPTION>	</CAPTION>	Table captions
Changed	<CHANGED>	</CHANGED>	Mark changed text
Command name	<CMD>	</CMD>	Set command name
Definition	<DFN>	</DFN>	Define instance of a term
Figure	<FIG>	</FIG>	Embed a figure and acts as a paragraph
Footnote	<FOOTNOTE>	</FOOTNOTE>	For additional information on some point
HTML+	<HTMLPLUS>	</HTMLPLUS>	Define HTML+ document
Image	<IMAGE>	</IMAGE>	Embed an image
Line break	<L>	no closing tag	Make explicit line break
Literal	<LIT>	</LIT>	Embed literal texts
Margin	<MARGIN>	</MARGIN>	Mark with margin attention label
Math	$	$	Embed mathematical equations
NextID	<NEXTID>	no closing tag	Generate identifier for anchor points

Table 19 (Page 2 of 2). HTML+ New Elements

Name	Opening tag	Closing tag	Description
Note	<NOTE>	</NOTE>	Bring attention to a point
Over	<OVER>	no closing tag	Divide math boxes into numerator and denominator
Person	<PERSON>	</PERSON>	Embed proper names
Quotation	<QUOTE>	</QUOTE>	Quote portions of text
Render	<RENDER>	no closing tag	Tell browser how to render unknown tags
Strike through	< S >	< / S >	Strikes a line through the font
Subscript	< S U B >	< / S U B >	Subscript text
Superscript	< S U P >	< / S U P >	Superscript text
Table	< T A B L E >	< / T A B L E >	Define a table
Table cell data	< T D >	no closing tag	Define table cell data
Table header(s)	< T H >	no closing tag	Define table's row header(s)
Table row	< T R >	no closing tag	Define table's row data

For more information on HTML+, see the following URL:

`http://www.yahoo.com/Computers/World_Wide_Web/HTML/HTML_3_0/`

Changes in the <BODY> tag:

- Backgrounds
- Colors

Table 20. <BODY> Tag Variables.

Variable	Description
BACKGROUND=	Points to a .gif image to use for the document background.
BGCOLOR=	Specifies the background color of the document, using a six-digit hexadecimal string. The string represents a mixture of red, green, and blue colors. (The first pair of digits represents red, the second pair green, and the third pair blue). A string in the form ″#000000″ generates a black background. You can view different color mixtures using the Color Palette editor in OS/2 Warp. This tag overrides the default settings in WebExplorer.
TEXT=	Specifies the color of the document text, using a six-digit hexadecimal string. For example, the string ″#CACA03″ generates yellow text. This tag overrides the default settings in WebExplorer.
LINK=	Specifies the color of links in the document, using a six-digit hexadecimal string. For example, the string ″#FF0000″ displays red document links. This tag overrides the default settings in WebExplorer.
VLINK=	Specifies the color of visited links in the document, using a six-digit hexadecimal string. This tag overrides the default settings in WebExplorer.

To use the <BODY> tag variables, you must put them inside the <BODY> tag. For example:

```
<BODY BACKGROUND=filename>
```

or

```
<BODY BGCOLOR=bgcolor TEXT=txtcolor LINK=lkcolor VLINK=vlkcolor>
```

filename The file name of the gif file to be used as your background.

bgcolor The six-digit hexadecimal string of the color you choose for your background.

txtcolor The six-digit hexadecimal string of the color you choose for your text in the document.

lkcolor The six-digit hexadecimal string of the color you choose for your links in the document.

vlkcolor The six-digit hexadecimal string of the color you choose for your visited links in the document.

4.1.3.1 Large Documents

HTML+ provides a way to split large documents over several servers to improve performance. A sequence of the document parts to be retrieved is established based on the assumptions that these documents are read from the beginning through the end; this sequence is known as a path.

In HTML+, the path can be declared at the beginning of the document, using the <LINK> tag. This tag can also be used to define glossary menu items suited for documents with many technical or unfamiliar terms or to provide a search field in every document page where readers can search by keywords. The tendency is to split a book into separate sessions as follows:

- Cover
- About the author
- Copyright
- Table of contents
- Foreword
- Preface
- Acknowledgement
- Chapters
- Appendix
- Bibliography
- Glossary
- Index

Each one of these sessions should be put into a separate HTML+ document. The table of contents should include hypertext links to other parts of the book.

4.1.3.2 Tables

Support for tables is one of the main enhancements of HTML+ over HTML. In this section, we will see how to create tables with captions, headers and data. Here we list some examples of applications. The table is declared using the <TABLE> tag; the caption is declared using the <CAPTION> tag. Table rows are declared using the tag <TR>, while the tags <TH> and <TD> define, respectively, table headers and table data. The BORDER attribute tells the browser to draw lines enclosing each table cell. Text in each cell is centered by default. A simple HTML+ table coding would look like the following:

```
<TABLE BORDER>
<CAPTION>Simple Table</CAPTION>
<TH>Col 1<TH>Col 2 <TH>Col 3 <TR>
<TD>1,1 <TD>1,2 <TD>1,3 <TR>
<TD>2,1 <TD>2,2 <TD>2,3 <TR>
<TD>3,1 <TD>3,2 <TD>3,3
</TABLE>
```

Figure 107 shows how a browser supporting HTML+ displays the table. In this example, we use the Arena browser for AIX.

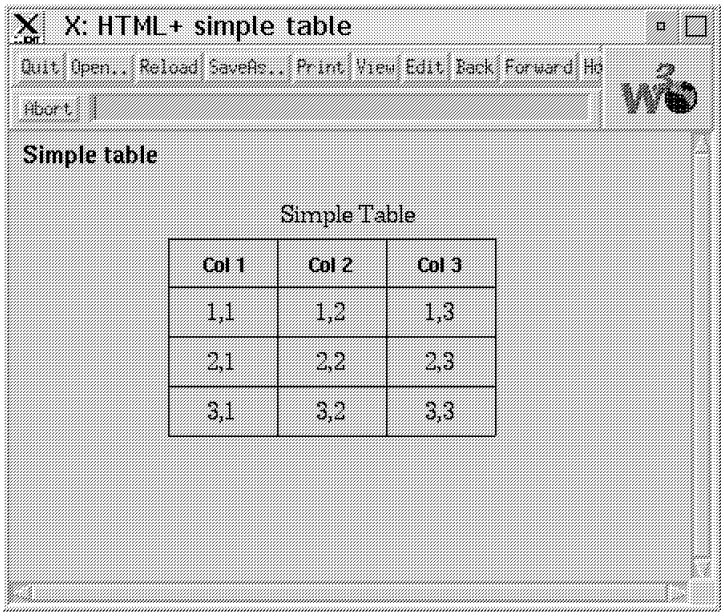


Figure 107. HTML+ Table

HTML+ supports the creation of more complex tables using other options, such as ROWSPAN or COLSPAN, that can define wider or higher cells in the table. The following example shows how to use these parameters:

```
<TABLE BORDER>
<CAPTION>Complex Table</CAPTION>
<TH>Col 1<TH>Col 2 <TH>Col 3 <TR>
<TD COLSPAN=2>1,1 and 1,2 <TD >1,3 <TR>
<TD>2,1 <TD>2,2 <TD ROWSPAN=2>2,3 and 3,3 <TR>
<TD>3,1<TD>3,2
</TABLE>
```

Figure 108 shows how the AIX Arena browser displays the table.

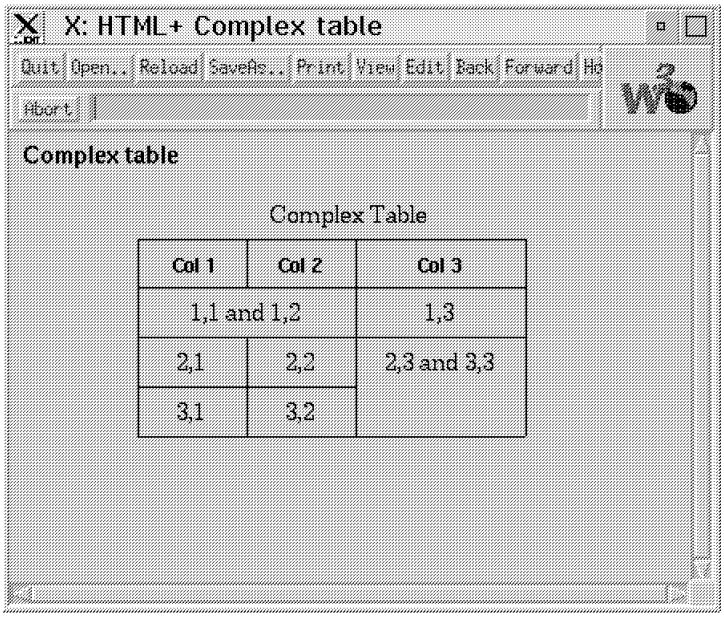


Figure 108. HTML+ Table

4.1.3.3 Mathematical Formulas and Equations

HTML+ supports the definition of mathematical formulas and equations. This is done by using the new <MATH> tags. The following example shows an HTML+ file that defines a few simple mathematical expressions:

```
<HTML>
<HEAD>
<TITLE>
HTML+ mathematical symbols
</TITLE>
<BODY>
<h2>
Mathematical symbols
<h2>
<h3>Equation</h3>
<MATH>
(a+b)<SUP>2</SUP> = a<SUP>2</SUP> +2 a b + b<SUP>2</SUP>
</MATH>
<h3>Equation</h3>
<MATH>
<BOX>(a<SUP>2</SUP> - b<SUP>2</SUP>)(a - b)<OVER>(a - b)<SUP>2</SUP></BOX>
 = (a + b)</MATH>
<h3>Equation</h3>
<MATH>
F<SUB>x</SUB> = m  <BOX>d<SUP>2</SUP>s<SUB>x</SUB>
<OVER>d t <SUP>2</SUP></BOX>
 = m <BOX> d<SUP>2</SUP> (s cos(&alpha;))<OVER> d t <SUP>2</SUP></BOX>
</MATH>
</BODY>
</HTML>
```

Figure 109 shows how Arena displays the table.

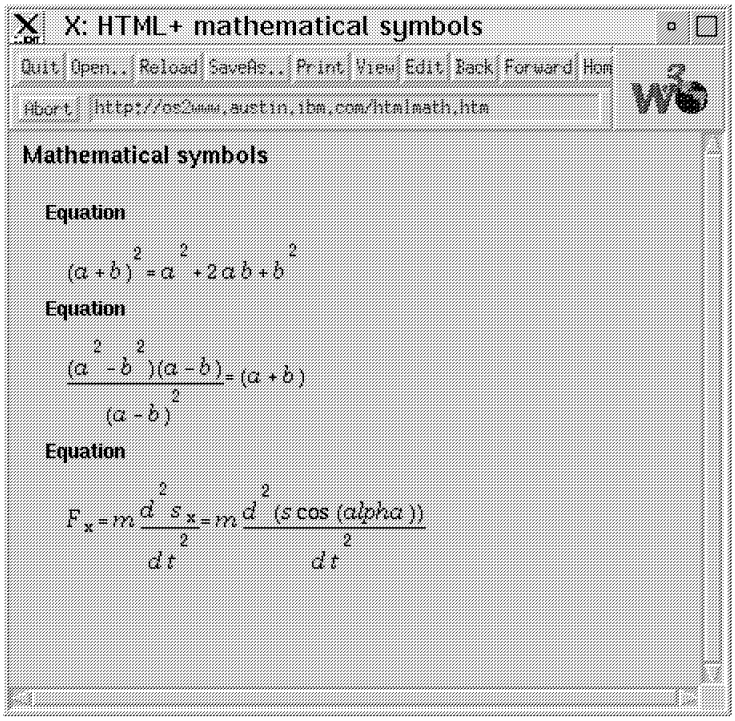


Figure 109. HTML+ Mathematical Expressions

4.1.4 HTML Special Symbols

As we have seen in the previous paragraphs, the symbols < (less than), > (greater than), & (ampersand), and ″ (double quote) are used to indicate tags in HTML language. If we want to show any of these symbols on the screen, we can′t just type them into the HTML source; the Web browser would attempt to interpret them as HTML tags.

To solve this, the following special commands have been defined to represent these symbols on the screen of a Web browser:

< is shown by the browser as <

> is shown by the browser as >

& is shown by the browser as &

" is shown by the browser as ″

HTML also supports extended characters. They are represented using symbols starting with the & character, as for example:

- é is written é
- ñ is written ñ
- ö is written ö
- ç is written ç

The following is an example of an HTML document written using special characters:

```
&lt;TITLE&gt;This is a title
&lt;/TITLE&gt; <P>
&lt;UL&gt; <P>
&lt;LI&gt;E acute: &eacute; <P>
&lt;LI&gt;C cedille: &ccedil; <P>
&lt;/UL&gt;
```

Figure 110 shows how this file is displayed by a Web browser.

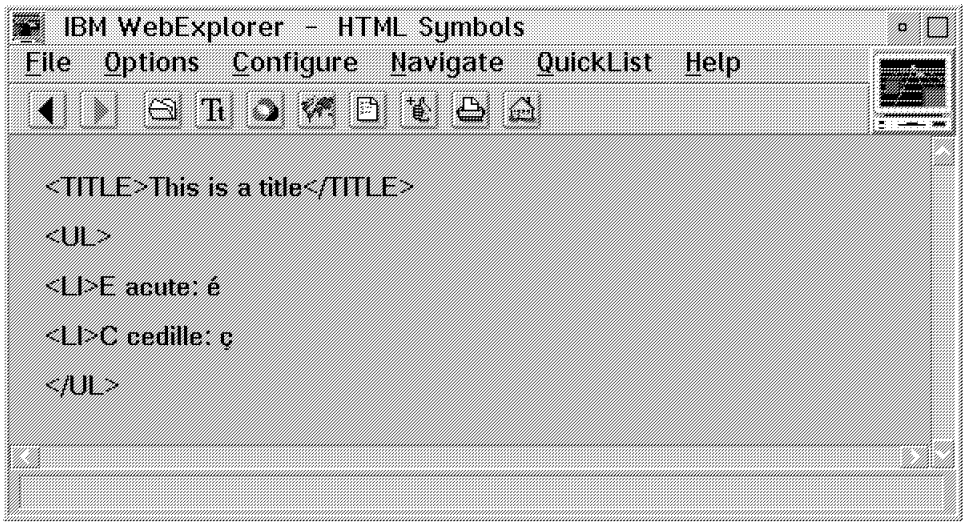


Figure 110. HTML Symbols

The list of extended character symbols can be found on an online HTML specification, such as the one at the following URL:

```
http://www.ucc.ie/info/net/html/
```

4.1.5 HTML Editors and Tools

All the examples and explanations in this chapter were based on the assumption that HTML documents were written using normal text editors. We showed parts of HTML document source, and separately, we showed how those documents were displayed by Web browser. This two-step process could be avoided using HTML editors.

4.1.5.1 IBM Electronic Publishing Edition for OS/2

The past several years have seen dramatic growth in the use of the Internet as a medium for electronic publishing. With IBM Electronic Publishing Edition for OS/2, documents can be created and served to internal corporate networks and to Hypertext Markup Language (HTML) browsers connected to the (WWW). And by utilizing BookManager READ, these same documents can be viewed by readers on multiple platforms who are not connected to an Internet Protocol Network.

Compared to the use of standard HTML and GIF files in other WWW libraries, IBM Electronic Publishing Edition for OS/2 offers significant advantages:

- BookManager format books are dynamically converted to HTML on demand.

- Each electronic book is a single readily portable and self-contained file, reducing the need to manage many separate HTML and GIF files.
- The BookManager book format allows much more content (up to 10 times more) to be stored on the same amount of disk space.
- A single server can serve books and bookshelves from its own storage or from multiple remote file systems. The actual location is not part of the Universal Resource Locator (URL) of the document and is transparent to the reader.
- Many document elements are supported beyond those directly supported in HTML, such as complex tables.
- Readers can use fuzzy and morphological full-text searching across entire documents and bookshelves not just the currently loaded HTML file.
- Navigation within documents is easier via a button bar with intuitive icons.

IBM Electronic Publishing Edition for OS/2 comes with everything needed to create and distribute documents on the WWW:

- IBM BookManager BUILD/2 Version 2.0 for building books from popular word processors (Microsoft Word, WordPerfect, AmiPro, and FrameMaker) files.
- IBM BookManager BUILD SGML for OS/2 Version 2.0 for building books from documents authored in Standard Generalized Markup Language (SGML).
- Language Dictionaries for building your books in multiple national languages.
- IBM BookManager BookServer for World Wide Web for OS/2 Version 2.0 for serving your books across the WWW.

Further information about IBM Electronic Publishing can be obtained at the URL `http://booksrv2.raleigh.ibm.com/`.

4.1.5.2 IBM HyperWise

This is an authoring tool that allows you to format and link text and graphics using drag and drop of OS/2 for HTML, GML and IPF.

HyperWise is a productivity tool for application and title developers. HyperWise enables What You See Is What You Get (WYSIWYG) authoring of hypertext on-line information and application help for OS/2 and Microsoft Windows.

With HyperWise, developers can use simple drag-and-drop techniques to link text, audio, video, and graphics. Developers can link to audio (.WAV and .MID), video (.AVI), and animation (.FLC and .FLI) extension files supported in WARP.

HyperWise Version 2.0, a replacement for Version 1.0, provides more editing features, enhances developer support for moving Windows help to OS/2, and supports World Wide Web browsers on the Internet. HyperWise 2.0 also helps you save time and resources; author the text once and read it on_ OS/2, Windows 3.1, and the Internet. Additional features of HyperWise 2.0 also make it easy for education specialists to create interactive tutorials for OS/2 applications.

The Information Presentation Facility (IPF) for Microsoft Windows is still packaged with HyperWise 2.0, so the same information compiled for OS/2 IPF is viewable on Windows. This single sourcing increases productivity and enables developers to use OS/2 for their development platform, regardless of the platform on which their applications run.

HyperWise 2.0 continues to require only limited disk space to store output. When HyperWise exports a readable format, it compresses text and graphics up to 80%.

Further information about IBM HyperWise can be obtained at the URL http://direct.boulder.ibm.com/us/desktop/appdev/p52c.htm.

4.1.5.3 HTML Editors

HTML editors are designed to get as close as possible to a what you see is what you get (WYSIWYG) approach. HTML editors usually have a menu from which markup tags can be selected and put into the text. For every tag there is a template that starts with the tag itself and contains information on the parameter and the syntax of the subject tag. List items are automatically indented as they are inserted. Every time a new HTML file is being created, the editor shows a template with all the tags that should always be included in HTML documents.

Here are a few of the more popular HTML editors running on various platforms and a URL where you can find more information about each editor:

- UNIX Platforms
 - ASHE

 ftp://ftp.cs.rpi.edu/pub/puninj/ASHE/README.html
 - tkHTML

 http://weber.u.washington.edu/~ roland/tkHTML/tkHTML.html
 - HoTMetaL

 http://www.sq.com/
 - Cyberleaf

 http://www.ileaf.com/ip.html
- OS/2
 - HTML Wizard

 ftp://ftp.cdrom.com/pub/os2/editors/htmlwiz.zip
 - HomePage Publisher

 ftp://ftp.apical.com/pub/HPP
- Windows
 - CU HTML for Word 6.0

 http://www.cuhk.hk/csc/cu_html/cu_html.htm
 - GT HTML for Word 6.0

 http://www.gatech.edu/word_html/release.htm
 - HoTMetaL

 http://www.sq.com/
 - HTML Author for Word 6.0

 http://www.salford.ac.uk/iti/gsc/htmlauth/summary.html
- Macintosh
 - html-helper-mode

 http://www.santafe.edu/~ nelson/tools/documentation.html

- NeXTStep
 - Pages

 http://www.pages.com/

More recent editors on all platforms can be found at the URL http://www.shareware.com.

4.1.5.4 HTML Tools

HTML editors are not the only software that has been developed to support the creation of HTML documents and WWW publishing; some HTML error checkers are also available on the Internet.

HTML Validation Service, for example, is available at the following URL:

http://www.hal.com/%7Econnolly/html-test/service/validation-form.html

The WWW page itself is the application user interface. It provides an entry field where the URL of the document to be checked must be entered and a validation level has to be specified. In case of heavy use of this tool, local installation is suggested. Figure 111 shows how this page looks when displayed by a Web browser.

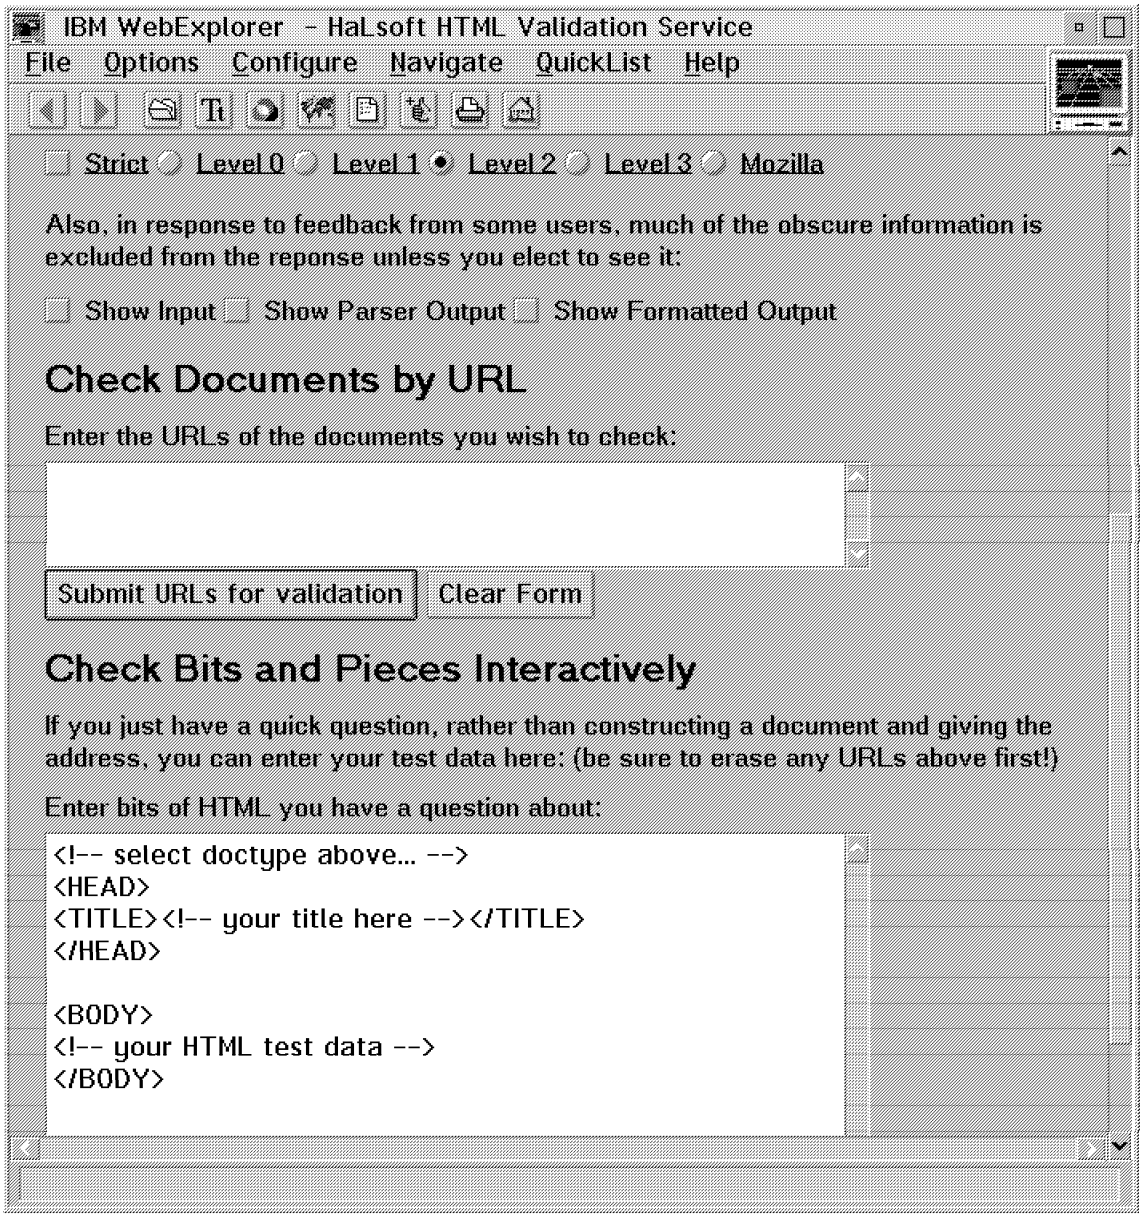


Figure 111. HaL HTML Validation Service

Another interesting tool can be found at the following URL:

`http://wsk.eit.com/wsk/dist/doc/admin/webtest/verify_links.html`

This tool starts the link verification at a given URL and traverses all the pointed links producing a report.

A tool called Weblint is also available by anonymous FTP at the following location:

`ftp://ftp.khoros.unm.edu/pub/perl/www`

For more information on this tool, its WWW page is located at URL:

`http://www.khoros.unm.edu/staff/neilb/weblint.html`

A syntax checker for HTML Versions 2.0 and 3.0 that includes other HTML utilities is available at the following URL:

http://uts.cc.utexas.edu/~churchh/htmlchek.html

4.1.6 Extensions to HTML

Some Web browsers can exploit some additional browsing capabilities given by an extended set of HTML tags and attributes. This is the case of the Netscape Web browser. The Netscape browser interprets more tags and commands than the standard ones defined for HTML. These are nonstandard HTML commands; they are disregarded by the other Web browsers.

Some of the additional features are:

- Customized message for ISINDEX search fields
- Additional parameters to HR (horizontal rule) HTML tag to specify line length
- Additional unordered list parameter to specify bullet shape
- Additional ordered list parameter to specify number or letters ordering
- Additional image alignment options
- No break tag, <NOBR>
- Word break tag, <WBR>
- Font size tag,
- Base font size tag, <BASEFONT SIZE = value>
- Center text tag, <CENTER>

A detailed reference of the Netscape extensions to HTML can be found at the following URL:

http://home.mcom.com/services_docs/html-extensions.html

4.2 Images

Images are an important part of World Wide Web documents. In this section, we analyze some details of the format of images to be embedded in HTML documents, their characteristics and related tools.

4.2.1 HTML Image Files

Graphic Web browsers can display HTML documents with in-line images. Generally, browsers can support multiple image formats; there is not an official image standard for Web publishing. However, the most commonly used format is GIF. If you create your images in GIF format you can be reasonably assured that your images will be viewable by most browsers.

Here are some of the graphic formats that you may encounter on the Web.

4.2.1.1 GIF

Graphics Interchange Format (GIF) is a commercial format still widely used on the Web. It was developed by CompuServe in 1987, and then revised in 1989 (GIF89) for additional capabilities.

The Graphical Interchange Format allows one-bit transparency so that images can be converted to transparent images. The GIF format uses a color table of 256 colors. The table can either be global, used by all the GIF images, or local.

When it is used locally, it is used by the image immediately following the table, and it supersedes the global table.

4.2.1.2 JPEG

Another graphic format used in Web documents is the Joint Photographic Expert Group standard (JPEG). JPEG compression methods can greatly reduce the image file size. A JPEG photographic image can produce a file 10 times smaller than the equivalent GIF. The standard is not recommended for images that have already been reduced to a 256-color palette.

4.2.1.3 PostScript

PostScript standard is a proprietary format whose usage is free. It is the world′s most popular standard to present text and graphics in a device-independent format. PostScript images can be displayed by tools, such as Ghostscript, available on AIX, OS/2, Windows, and Macintosh platforms, and Ghostview, available on AIX and Windows platforms. Applications that display PostScript files are also freely available on the Internet. The big advantage of PostScript is that, since it is such a common printer language, almost all applications can produce it. The drawback is its extensive use of macros, sometimes not optimized by the application producing the PostScript files. This causes these files to be very large.

4.2.2 PDF (Portable Document Format)

This format is a proprietary format from Adobe Systems Incorporated that allows you to create multiple-page documents and create internal links on them, having all the advantages of the PostScript as well. Readers for this format can be found for OS/2, Windows 95, Windows NT, Windows 3.1, Macintosh, SPARC Sun OS, SPARC Solaris, HP-UX, IBM AIX and Silicon Graphics IRIX. All download readers are at:

`http://www.adobe.com/acrobat.`

4.2.3 Transparent Images

Transparent images are images whose background color matches the color of the browser′s background, giving the impression that they are floating on top of the document.

Some Web browsers have configuration options that allow the users to customize the colors; so the transparency effect can′t be obtained by giving the image background a certain color because a user′s settings of the browser are various and unpredictable. These images really must have a transparent background. In Figure 112 we show a Web browser page displaying a normal and a transparent GIF image.

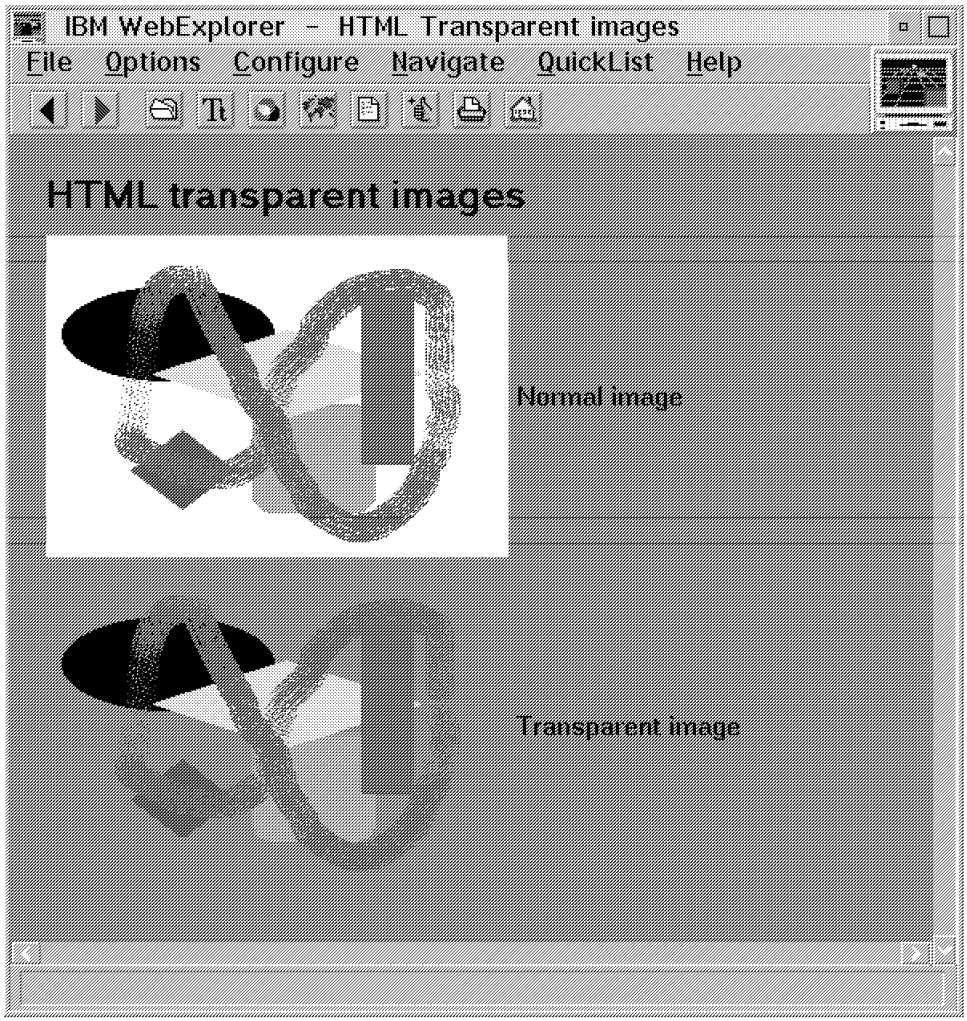


Figure 112. Images - Transparent GIF

Here, we describe the steps of the process to be followed to transform a normal GIF image into a transparent image.

The GIF image must be generated, captured from the screen or downloaded from any online image archive. There must be only one color in the image background, and this color shouldn't have been used anywhere else in the image because all the parts of the image painted with that color will become transparent.

4.2.3.1 Making Transparent Images

The only image format that supports the transparency feature is the GIF89a. If the image to be processed is GIF87a, it must be converted. This can be done by a tool called giftrans, available by anonymous FTP from the following URLs:

```
ftp://pascal.ibp.fr/pub2/www/tools/
ftp://lune.csc.liv.ac.uk/hpux/X11/Graphics/giftrans-1.11.1/
ftp://ftp.sunet.se/pub/www/utilities/www-tools_uni-karlsruhe/
ftp://sgml1.ex.ac.uk/pub/WWW/msdos/editors/
```

Giftrans can convert GIF87a to GIF89a transparent in one step. The program is run by typing the following command:

```
giftrans -t index GIF87afn > GIF89afn
```

where:

GIF87afn	filename of the input GIF87a image file
GIF89afn	filename of the output GIF89a image file
index	hexadecimal RGB triple of the color to be made transparent

Some useful image converters are also available on the Web; they can be found at the following URLs:

```
http://www.vrl.com/Imaging/convert.html
http://www.vrl.com/Imaging/transparent.html
```

The first one is an on-the-fly image format converter; its user interface is the Web Page itself. A number of options can be selected for the conversion, and the tool can retrieve our local image to process it. A drawback to using tools such as these is that the tool needs to be able to retrieve your image in order to convert it. If your system is located inside of a firewall, the tool will not be able to retrieve your image because the firewall will block its access to your system. The only way around this is to ask your system administrator to put your image on your organization′s external Web server. This will allow the converter tool to retrieve your image and convert it as desired. If your system is not inside of a firewall, you need to make your image available on a Web server so that the tool can retrieve it for the conversion. Ask your Service Provider if they can help you out by placing your image on their server.

4.3 Other Resources (Audio and Video)

Other resources, such as video and audio clips, can easily be included in your HTML documents. In fact, anything that is not text or an image can be included using this simple procedure. To include these kinds of resources, you simply put a hyperlink to the resource in your document. For example, if you wanted to add an audio clip into your document, you would simply include a hyperlink such as the following in your document. The URL in the hyperlink points to the address of the audio file that should be played when the hyperlink is selected.

```
<a href=http://myserv/myvoice.wav>Click here to hear my voice</a>
```

The file myvoice.wav, which is served by the Web server named myserv, is a data file that contains an audio clip of your voice that has been digitized and saved in the file using one of the standard audio formats. When the reader selects the hyperlink, the browser will request the file specified in the URL from the server also specified in the URL. When the server transfers the file back to the browser, the browser will determine the MIME-type of the returned data file and call the appropriate external viewer to play the audio clip for the reader. The process is exactly the same for any other non-text or image resource. You simply:

1. Create the resource (data file).
2. Place it on a Web server.
3. Hyperlink to it in your document.
4. Let the reader worry about configuring their browser to call an appropriate viewer on their platform to handle the resource file. Of course, it would be polite if you included information in your document on the nature and format of the resource so the reader can easily configure their viewer.

4.4 HTML Converters

The Hypertext Markup Language is the standard language for creating documents for the World Wide Web. Every document published on the Web should conform to this standard. There are cases where it might be necessary to author documents in other languages and/or systems and then convert the document to HTML. These include:

- Some authors might not know how to write in HTML.
- You may have previously written documents that you want to make available.
- You may need to develop the document in a specific format. For example, you may want to also publish a hardcopy of the document, and your publisher may require the document in a specific format.

Regardless of the reason, documents created in formats other than HTML can, in most cases, be easily converted using one of several format conversion tools or filters. The output of these tools is seldom perfect HTML format. However, the output is usually close and generally only requires a little cleanup or the addition of the hyperlinks. Therefore, a knowledge of HTML is still required in order to modify the document for distribution. In this section, we describe a few of the more popular HTML converters currently available. Information on lots of other converters can be found at the following URL:

```
http://union.ncsa.uiuc.edu/HyperNews/get/www/html/converters.html
```

The following sections cover conversion from:

- BookMaster to HTML
- FrameMaker to HTML
- Interleaf to HTML

4.4.1 BookMaster to HTML

The conversion from BookMaster to HTML is done by a program called BookMaster Utility; the executable file is called bk2html, which is written in C++ on OS/2 2.1 by Martin Tasker of Imperial College, London. IBM BookMaster is a markup language used to write documents. BookMaster tags begin with a colon and end with a dot. Their names are sequences of alphanumeric characters and can have attributes to be specified inside the tag delimiters (the colon and dot). All colons that are not followed by a blank are treated as beginnings of a tag. Large BookMaster documents are generally split into several modules; a main file embeds all the modules with the .im macro. The main BookMaster markup tags are:

:p. Begin paragraph

:h1-20. Define up to twenty levels of heading

:hp1-9. Define up to nine highlighting levels

:ul. Define an unordered list

:ol. Define an ordered list

:dl. Define a definition list

:li. Define a list item

:cit. Define italicized citations

:index. Build index

:i1. Create index entry

:toc. Build table of contents

:fig. Begin figure

:table. Define a table

Detailed information about IBM BookMaster can be found in the *IBM BookMaster User's Guide 4.0*.

bk2html runs under OS/2 and AIX. It is invoked by typing the following on the command line:

```
bk2html <options> fn<.ext>

where:

        fn          filename of input file to be processed
        ext         extension of input file (default .SCR)

Options:

        -f format
                    select output format: (default html)
                                 html      format for HTML WWW browser
                                 latex     format for LaTex processing

        -m mainfn
                    specify main Table Of Content file (default MAINFN.TOC)

        -od outdir (default current directory)
                    specify output directory
```

bk2html converts the input source BookMaster files into HTML language according to HTML, March 1994, CERN specifications; output files will have the .HTML extension in UNIX and the .HTM extension in OS/2.

bk2html generates one output file for each processed input file and for each file embedded by the input file using the .im macro. Whenever this .im macro is found, bk2html generates a HyperText anchor of the type <HREF=″embedded file″> in the output file that points to the first heading of the embedded file. Also, links to referenced headings are supported. A BookMaster reference looks like the following:

```
:h1 id=alpha.Alpha
```

It cross references to:

```
:hdref refid=alpha.
```

It is converted in the reference:

```
<h1><A NAME=alpha>Alpha</A></H1>
```

The cross reference link will be:

```
<A HREF=#alpha>Alpha</A>
```

Here is an example of a simple BookMaster file conversion to HTML. The source BookMaster file is named bktohtml.script.

```
:h1 id=title.BookMaster to HTML Conversion
This sample script file will include the following marks:
:sl.
:li.Heading (level 1 and 2)
:li.Unordered lists, see :hdref refid=lists.
:li.Cross reference
:esl.
:h2 id=lists.Lists
There are four kinds of lists:
:ol.
:li.Ordered lists
:li.Unordered lists
:li.Definition lists
:li.Simple lists
:eol.
```

Figure 113 shows how this file is formatted by BookMaster.

BookMaster to HTML conversion

This sample script file will include the following marks:

- Heading (level 1 and 2)
- Unordered lists, see ″Lists″
- Cross reference

Lists

There are four kinds of lists:

1. Ordered lists
2. Unordered lists
3. Definition lists
4. Simple lists

Figure 113. BookMaster Formatting

The file was converted by entering the following syntax from a UNIX command prompt:

```
bk2html -f html bktohtml.script
```

The output file, bktohtml.html, is as follows:

```
<!-- output file generated by BM Utilities -->
<html>
<head>
<body>


<hr>
<h1><a name="title">BookMaster to HTML conversion</a></h1>
This sample script file will include the following marks:
<menu>
<li>Heading (level 1 and 2)
<li>Unordered lists, see <a href="#lists">Lists</a>
<li>Cross reference
</menu>

<h2><a name="lists">Lists</a></h2>
There are four kinds of lists:
<ol>
<li>Ordered lists
<li>Unordered lists
<li>Definition lists
<li>Simple lists
</ol>
```

Figure 114 shows how this file is formatted by a Web browser.

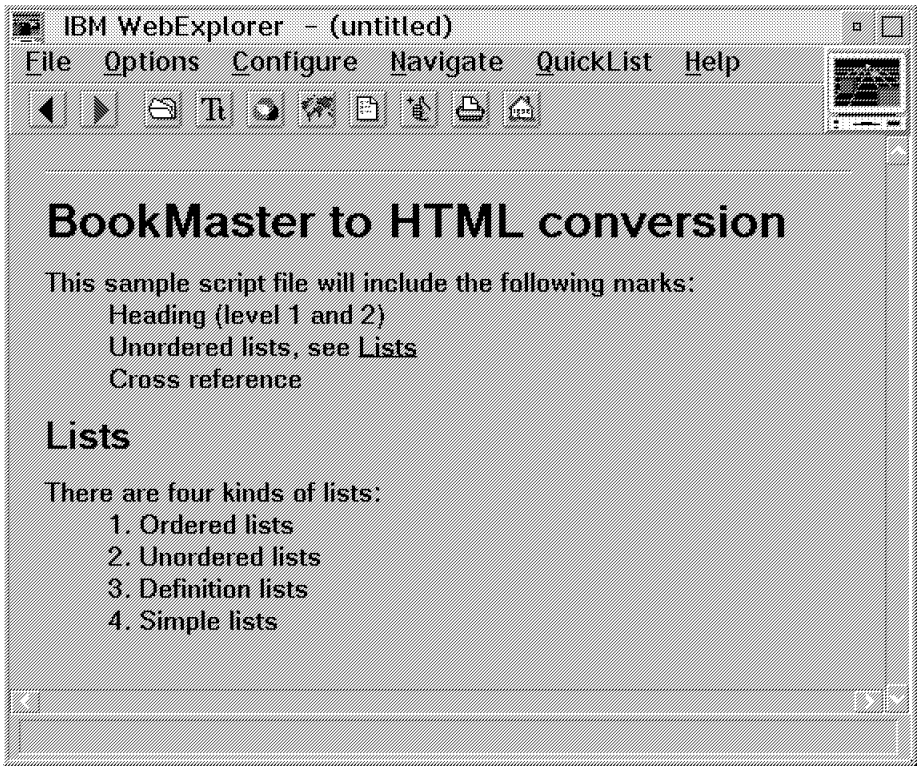


Figure 114. BookMaster to HTML - Converted Document

Here is a list of enhancements that the author is planning to make to the program:

- Multiple input directories support
- OS/2 or Windows help support
- Reference to other books in the same library support
- Table support
- Mathematical formulas support

bk2html can be found at the following London Imperial College URL:

```
http://rankine.cv.ic.ac.uk/
```

4.4.2 FrameMaker to HTML

The conversion from FrameMaker to HTML is done by two different programs.

- fm2html - for FrameMaker Version 3.0 documents
- WebMaker - for FrameMaker Version 4.0 documents

This program can convert FrameMaker documents and books and supports conversion of figures, mathematical formulas and tables.

FrameMaker documents are logically structured and contain specification of contents and layout. FrameMaker documents can be divided into the following four main sections:

- Structure specification
- Tables and frames specification
- Page layout information
- Text paragraph with reference to other paragraphs

Before being converted to HTML, FrameMaker files have to be turned into the FrameMaker Interchange Format (MIF) by calling the FrameMaker program fmbatch. fm2html converts from MIF format to HTML format. During the generation of the MIF file, figures are extracted and put into separate files, and a table of contents is generated. Conversion of FrameMaker books follows the same process of single files conversion.

MIF files contain a lot of information regarding the FrameMaker document. The part of this information needed by HTML is converted; the rest is ignored. In HTML, for instance, page numbers do not have meaning since HTML documents are seen entirely in a flow. Every FrameMaker reference to a page number is ignored by the converter. FrameMaker uses hypertext links. All these links, except for the ones referencing page numbers, are converted into HTML anchors. FrameMaker footnotes and references are also converted into HTML anchors.

FrameMaker can include figures in different formats. During the conversion process these figures are converted into GIF format; that is, the image format recognized by all the graphical Web browsers.

The current version of HTML does not include support for tables and mathematical expressions; the only way to include them into HTML is to transform them into figures before using the converter.

Further information about fm2html can be found at the following URL:

```
http://www.w3.org/pub/WWW/Tools/fm2html.html
```

Further information about WebMaker can be found at the following URL:

```
http://www.cern.ch/WebMaker/
```

4.4.3 Interleaf to HTML

The conversion from Interleaf to HTML is done by a program called il2html. Interleaf for Motif is a software product for document creation, composition and assembly that supports hypertext links, embedded figures, tables, and mathematical equations.

Before being converted, Interleaf documents must be saved in Interleaf ASCII format. This can be done by Interleaf itself by choosing the option:

Save --->ASCII - Forced

il2html is invoked by typing the following on the command line:

```
il2html filename.doc > filename.html
```

Where:

```
filename.doc      filename of input file to be processed
Filename.html     filename of output file
```

Text conversion is completely automatic; for graphics, some hand work is still required. The filter just includes an empty image reference:

```
<IMG SRC=" ">
```

The following is the recommended step-by-step process to be followed for graphics creation:

1. Start Interleaf.
2. When the main window appears, click the right mouse button to bring up the controls.
3. Grab the image to be converted by clicking on **Grab** and moving the mouse to draw a box around the image.
4. Save the picture by clicking on **Save**.
5. Use the GIF format and the full color option; save in a file with the .gif extension.
6. Quit.

Once the image is created this way, the HTML file must be modified to insert the image. The SRC= field must be filled with the path and file name of the image.

Further information about il2html can be found at the following URL:

```
http://18.23.0.23/pub/WWW/Tools/il2html.html
```

An Interleaf to HTML converter has been developed by Interleaf, too; its name is iam2html. Once Interleaf files have been saved to Interleaf ASCII format, conversion can be done by typing:

```
iam2html filename
```

This will produce an output file named filename.html.

Information on this product can be found at the WWW Interleaf page at the following URL:

```
http://www.ileaf.com
```

Sometimes it can be more convenient to convert Interleaf files to FrameMaker and then to HTML. Conversion from Interleaf 4.0 to FrameMaker can be done by Filtrix, a commercially licensed package developed by Blueberry Software. Interleaf documents must be saved in Interleaf 4.0 ASCII format. Once Filtrix is started, the source directory must be changed to the directory where the files to be converted are stored. Every file in this directory will be listed; the files to be converted can be selected with a mouse click and their output name must be specified. A .mif extension is recommended. Files are now ready to be processed by the FrameMaker to HTML converter.

Interleaf has a commercial product called Cyberleaf that also does Interleaf to HTML conversions. More Information about Cyberleaf can be obtained at the following URL:

`http://www.ileaf.com/ip.html`

4.4.4 Other HTML Converters

The following is a partial list of some other popular HTML converters available and the locations on the Internet where further documentation can be found.

`http://union.ncsa.uiuc.edu/HyperNews/get/www/html/converters.html`

- Postscript to HTML

 `http://www.area.fi.cnr.it//area/ps2html.htm`

- Lotus Notes to HTML

 `http://tile.net/info/about.html`

- LaTex to HTML

 `http://cbl.leeds.ac.uk/nikos/tex2html/doc/latex2html/latex2html.html`

- PageMaker to HTML

 `http://www.bucknell.edu/bucknellian/dave/`

- PowerPoint to HTML

 `http://www.w3.org/hypertext/WWW/Tools/PPT.html`

- C++ to HTML

 `http://www.bauv.unibw-muenchen.de/graphics/projects/c++2html.html`

- Fortran to HTML

 `http://vscrna.cern.ch/floppy/contents.html`

4.5 CGI's Programming

In order to make a complete reference of the standard and create a background before doing such an analysis, we are making a technical approach first. After the CGI specifications, there are some examples and their analysis. In this way you can have a quick reference at the beginning and a practical one at the end.

CGI, which stands for Common Gateway Interface, is only a programming standard to communicate with the web server and the WWW with your program. The steps you have to follow to make a CGI program are: CGI programs have to be in a directory with executed permissions by the web server; if you have an IBM web server you already have 2 directories with those permissions: cgi-bin and admin-bin. If you want to create a new one, use the administration forms using the request routing option.

In other CERN-based servers you also have the cgi-bin directory.

1. The choice of the transference method.
2. The environment variable catch to know the transference method.
3. The catch of the ″QUERY_STRING″ environment variable if the selected method was GET.
4. The standard inputs are driven by the Web Server if POST method is implemented.
5. The standard outputs are overridden to the client (browser).
6. The standard output must have a header.
7. The standard input stays with special separators, the same as the QUERY_STRING variable.

4.5.1 The Choice of the Transference Method

CGI has different transference methods of interaction between the server and the client; the best known are GET and POST. These methods allow the programmer to take control of the data in an easy way.

To know the method that the client (browser) implemented for the data transference, the CGI program has to look in the REQUEST_METHOD environmental variable, in order to look at what type of decoding has to be used. So we already note that the client is the one who is going to choose the method. But how?

When we make a form using HTML we put the method that has to be used by the browser:

```
<FORM ACTION=″/cgi-bin/mycgi″ METHOD=″GET″>
```

You can use either GET or POST on the form.

4.5.2 Catching the REQUEST_METHOD Variable

As you can see CGI is a normal program with too little specifications. To get the method used by the client you only have to use the correct command to get those variables. Example in c:

```
char * method;
.
.
.
 method=getenv(″REQUEST_METHOD″);
 if (!strcmp(method, ″GET″)) /*The chosen method is GET*/
                                ....
 if (!strcmp(method,″POST″) /*The chosen method is POST*/
                                ....
```

If you are using other languages such as REXX, PERL, and VisualBasic, you′ll have to use the equivalent command.

4.5.3 Catching the QUERY_STRING Variable

You are going to do the same as you did in the last step:

```
char * information;
.
.
.
if (!strcmp(method,"GET"))    /*You only have to look for the
                                 QUERY_STRING if the method is GET*/
    information=getenv("QUERY_STRING");
```

The information has to be decoded (see step 5).

4.5.4 Standard Input on the POST Method

You are going to use the standard input stream instead of the QUERY_STRING if the POST method is implemented.

Important Note

If you use the GET method, the information that you send is part of the URL; if you use the POST method, the information that you send is *not* part of the URL and you can put it into a variable reading the standard input. For example: http://www.ibm.com/cgi-bin/cgiprogram?information=time+to+sleep uses the GET method, and http://www.ibm.com/cgi-bin/cgiprogram uses the POST method (if some information was send).

4.5.4.1 CONTENT_LENGTH

This variable gives you the number of bytes of the said content by the client. Knowing this variable allows you to open a standard output like a stream and directly read the quantity of bytes the client send to you.

```
FILE *f;
.
.
.
f=stdin;
if (feof(f)){ /*Something happened on the stdin and we can't read*/
    printf("Content-type: text/html\n\n");
    printf("An error ocurred when the server tried to get your \
            information");
    }
else
    {
     information=fread(f,atoi(getenv("CONTENT_LENGTH")));
    }
```

The next step you have to do to use the information is to decode it.

4.5.5 Standard Output

The sever will send all the standard outputs to the client, but you must tell the client the type of data you are sending before starting. The way to tell the client what the content is, is to make the first standard output with the following format:
Content-type: MIME TYPE

This line must be followed by a blank line (two new line characters) and the content you send. For example:

```
.
.
.
printf("Content-type: text/html\n\n");
printf("<HTML>\n<HEAD><Title>Succesfull transaction</Title>");
printf("<i>Your transaction was successfull.<p>");
printf("<A href=\"/\">Return to home</A>");
.
.
.
```

If you want to send an image you have to change the contents type to an image/gif for example. Look in the CD-ROM for CGI programs examples; you have animator source codes, text file writing programs, and UNIX mail senders programs. One of the most important things on CGI programming is to use the KISS (Keep It Short and Simple) philosophy. You normally won't need programs too large or complex.

4.5.6 Decode the Input

The input must be decoded to get the information you need. You can use a 2-string structure to get the information right where you need it.

The structure could be something like this:

```
typedef struct {
char variable [25 ] ;
char value [ 1024 ];
} decode;
```

Note that you are putting a limit in the amount of 1024 characters. If you are going to use it with a form that uses text area, we highly recommend you make this value for about 32 K or more. The information is coded this way:

1. If the method is GET, the information is part of the URL making the separation with an interrogation mark (?) between both. The part that you have to decode doesn't have this interrogation mark and is on the QUERY_STRING variable.
2. Every variable and its contents are separated by an ampersand (&) from each other. The last couple of variable values has no ampersand at the end of it.
3. The variable name is separated by an equal sign from their value (name=Roberto+Oku).

The first thing you have to do is to separate every variable from the others and then separate the name from their value.

On the CD-ROM you find a file named util.c that implements these features and two examples of queries: post_query.c and query.c that implement the catching of each variable. These files are freeware and you can also get them from:

`ftp://ftp.ncsa.uiuc.edu/Web/httpd/Unix/ncsa_httpd/cgi/cgi-src`

4.5.7 CGI Variables

This section covers other useful variables on the CGI standard that you should know.

4.5.7.1 SERVER_SOFTWARE

You will find the name and version of your server with the following format: name/version. The software makes the administration of your server or the administration for more than one server; this could help you know what the features you might need to change on the configuration files. This variable is not for any specific request, which mean all the requests are going to have it.

4.5.7.2 SERVER_NAME

This has the server name, DNS or IP address. It is the name that the server gives itself to make self-references requests or URL references. If you want to put a URL as a part of the output on your CGI you must use this environment variable instead of coding the name itself.

4.5.7.3 GATEWAY_INTERFACE

This contains the information about the CGI specifications you can use on the server. The list of variables and usage you are reading complies with the CGI Version 1.1. The format that is given is CGI/revision.

4.5.7.4 SERVER_PROTOCOL

It indicates what was the server protocol of the request. If you want to maintain only secure transactions you may respond only on those cases that have secure protocols such as SHTTP or SSL.

4.5.7.5 PATH_INFO

This is the extra path information that the client gives to perform the CGI program. This information has to be decoded for the server before the CGI programs perform its action. For example, in the db2www you can use something like this:

```
http://.../cgi-bin/db2www/report
```

The report parameter stays in the PATH_INFO variable.

4.5.7.6 PATH_TRANSLATED

This is a virtual to physical translation of the request.

4.5.7.7 SCRIPT_NAME

This is the virtual path name of the request, which is used to generate self-referenced links into the CGI program like output.

4.5.7.8 REMOTE_HOST

This is the name of the host who makes the request.

4.5.7.9 REMOTE_ADDR

This is the IP address from the requester host.

4.5.7.10 AUTH_TYPE

This is the protocol authentication method used to validate the user.

4.5.7.11 REMOTE_USER

This is the named user when authentication is set.

4.5.7.12 REMOTE_IDENT

If the HTTP server supports the RFC 931 authentication, this variable is set with the remote user name retrieved from the server. This is for logging purposes only.

4.5.7.13 CONTENT_TYPE

This variable contains the type of data transmitted on the transaction, if you are going to make a Form validation CGI you must check to make sure the contents are from a Form and not some other kind of data before decoding (see the following examples).

4.5.7.14 HTTP_ACCEPT

This gives you the MIME items the client can accept in response; you use it to know the browser's capabilities. Each item is formatted (type/subtype) and they are separated by commas.

4.5.7.15 HTTP_USER_AGENT

This gives you the software the client is using as a browser with the following format: software lybrary/version, allowing you to make multiple CGI responses based on the features of the browsers (such as Netscape Frames, multiparted/x-mixed-replace contents or Web explorer's <ANIMATION> tag).

4.5.8 Content Type considerations

As you'll see in the examples below where the method used is POST, you have to be careful with the type of information you are receiving in order to check the contents of the package received by the client.

The content for the post from a Form should be:
`application/x-www-form-urlencoded`.

4.5.9 Examples, Examples, Examples

Before checking the examples, we have to make certain the kind of considerations to implement the CGI. One of these has to be the language we are going to use.

A lot of people take script languages such as PERL or REXX, but this is not always the right answer to the problem.

It is faster to execute a program that has been compiled than a program that has to be interpreted, and, the greater the program, the greater the difference becomes. This is why we recommend you choose a languag such as C or C++ in order to make CGI programming.

If you don't want to write your code in C because you care about the transportability of the program (you may not want to compile the program in different machines), we will give you some hints for choosing an interpreter:

- The language has to be available on a wide variety of platforms.
- It has to be easy to understand and program.
- The interaction with the external environment has to be clean, transparent and powerful.

The language that we recommend you use, to do CGI programming if you want to program with an interpreted language, is REXX. You already have this language as the default interpreted language in OS/2, DOS, and VM Systems, and you can get UNIX versions (in Linux, AIX 3.2.5, HP UX 9.x Sun OS 4.1.3, Sun Solaris 2.4 and Silicon Graphics Irix 5.3) and even Amiga or Windows NT (from Microsoft Corp.). For more information on REXX and how to obtain the version, you need to access http://www2.hursley.ibm.com/rexx/.

4.5.9.1 NCSA Query

The following two programs intercept a Form contents and displays them into the browser with the variable = value pair. The query.c programs is used only for the GET method requests and post_query.c is used for the POST method. Both of them are on the CD-ROM and you can download them from: http://hoohoo.ncsa.uiuc.edu/cgi-forms.html.

```
#include <stdio.h>

#define LF 10
#define CR 13

void getword(char *word, char *line, char stop) {
    int x = 0,y;

    for(x=0;((line[x]) && (line[x] != stop));x++)
        word[x] = line[x];

    word[x] = '\0';
    if(line[x]) ++x;
    y=0;

    while(line[y++] = line[x++]);
}

char *makeword(char *line, char stop) {
    int x = 0,y;
    char *word = (char *) malloc(sizeof(char) * (strlen(line) + 1));

    for(x=0;((line[x]) && (line[x] != stop));x++)
        word[x] = line[x];

    word[x] = '\0';
    if(line[x]) ++x;
    y=0;

    while(line[y++] = line[x++]);
    return word;
}
```

```
char *fmakeword(FILE *f, char stop, int *cl) {
    int wsize;
    char *word;
    int ll;

    wsize = 102400;
    ll=0;
    word = (char *) malloc(sizeof(char) * (wsize + 1));

    while(1) {
        word[ll] = (char)fgetc(f);
        if(ll==wsize) {
            word[ll+1] = '\0';
            wsize+=102400;
            word = (char *)realloc(word,sizeof(char)*(wsize+1));
        }
        --(*cl);
        if((word[ll] == stop) || (feof(f)) || (!(*cl))) {
            if(word[ll] != stop) ll++;
            word[ll] = '\0';
            return word;
        }
        ++ll;
    }
}

char x2c(char *what) {
    register char digit;

    digit = (what[0] >= 'A' ? ((what[0] & 0xdf) - 'A')+10 : (what[0] - '0'));
    digit *= 16;
    digit += (what[1] >= 'A' ? ((what[1] & 0xdf) - 'A')+10 : (what[1] - '0'));
    return(digit);
}

void unescape_url(char *url) {
    register int x,y;

    for(x=0,y=0;url[y];++x,++y) {
        if((url[x] = url[y]) == '%') {
            url[x] = x2c(&url[y+1]);
            y+=2;
        }
    }
    url[x] = '\0';
}

void plustospace(char *str) {
    register int x;

    for(x=0;str[x];x++) if(str[x] == '+') str[x] = ' ';
}
```

Figure 115 (Part 2 of 2). util.c. Utilities for decoding from NCSA.

This file contains all the functions you need to decode the form you are posting, even if the method is GET or POST.

We have only mentioned some functions of the c file in order to focus on those that are important to us.

The getword function is important for decoding and obtaining the values from a string; you can note that the fmakeword function works the same way with the exception of reading from a file. The parameters used are the string we are looking for (an empty string) where we want to have the returned value which is going to be placed in a parameter and the line parameter which is going to be modified with the line without the word string. This is in order to find the strings that are separated for some special characters. In the URL encoded we are going to have 2 special cases, when we are separating the pairs of variable names and values and the one we are using to separate the name from the value. The first one is an ampersand (&) and the second one is an equal symbol (=).

The makeword function and the fmakeword function works in the same way, but they return the line value as the function′s return value instead of modifying the contents of the char pointer parameter.

Just keep in mind these functions; we are going to use them to process the Form′s information on the next two programs.

```
#include <stdio.h>
#ifndef NO_STDLIB_H
#include <stdlib.h>
#else
char *getenv();
#endif

typedef struct {
    char name[128];
    char val[128];
} entry;

void getword(char *word, char *line, char stop);
char x2c(char *what);
void unescape_url(char *url);
void plustospace(char *str);


main(int argc, char *argv[]) {
    entry entries[10000];
    register int x,m=0;
    char *cl;

    printf("Content-type: text/html%c%c",10,10);

    if(strcmp(getenv("REQUEST_METHOD"),"GET")) {
        printf("This script should be referenced with a METHOD of GET.\n");
        printf("If you don't understand this, see this ");
        printf("<A HREF=\"http://www.ncsa.uiuc.edu/SDG/
            Software/Mosaic/Docs/fill-out-forms/overview.html\">
            forms overview</A>.%c",10);
        exit(1);
    }
    cl = getenv("QUERY_STRING");
    if(cl == NULL) {
        printf("No query information to decode.\n");

        exit(1);
```

Figure 116 (Part 1 of 2). NCSA Example on the GET Method -query.c

```
    }
    for(x=0;cl[0] != '\0';x++) {
        m=x;
        getword(entries[x].val,cl,'&');
        plustospace(entries[x].val);
        unescape_url(entries[x].val);
        getword(entries[x].name,entries[x].val,'=');
    }

    printf("<H1>Query Results</H1>");
    printf("You submitted the following name/value pairs:<p>%c",10);
    printf("<ul>%c",10);

    for(x=0; x <= m; x++)
        printf("<li> <code>%s = %s</code>%c",entries[x].name,
                entries[x].val,10);
    printf("</ul>%c",10);
}
```

Figure 116 (Part 2 of 2). NCSA Example on the GET Method -query.c

This is the main example for the GET method. As you can see, as the first step on the main function, we are checking the value of the REQUEST_METHOD environment variable, then we look for the QUERY_STRING value and put it on the cl variable.

Important Notice on the Listing

If you know how to make a C program you already know that there's one line that was written in 3 lines due to the lack of space: the line says: `printf("<A HREF=\" . . . overview </A>. %c",10);` if you are copying the text you have to be careful with this.

Once we have the information to decode on the cl variable the decoding is made on the Form loop using the getword function; the word is kept on entries[x]val, the rest of the line remains on cl and the character we have for reference to make the partition is the ampersand (&) symbol. After this happens we have to take all the + symbols from the strings. These symbols represent spaces and have to be replaced (that's what the plus to space function does) by decoding the special unicode characters with the unscape_url function and putting the name of the variable in the name field and the value on the val one.

4.5.9.2 The post_query.c Example

The post_query example is very similar and has the same exit as query.c.

The steps to follow in the post_query are those that were described before: look for the environmental variables, read the standard input and decode it.

```
#include <stdio.h>
#ifndef NO_STDLIB_H
#include <stdlib.h>
#else
char *getenv();
#endif

#define MAX_ENTRIES 10000

typedef struct {
   char *name;
   char *val;
} entry;

char *makeword(char *line, char stop);
char *fmakeword(FILE *f, char stop, int *len);
char x2c(char *what);
void unescape_url(char *url);
void plustospace(char *str);

main(int argc, char *argv[]) {
   entry entries[MAX_ENTRIES];
   register int x,m=0;
   int cl;

   printf("Content-type: text/html%c%c",10,10);

   if(strcmp(getenv("REQUEST_METHOD"),"POST")) {
      printf("This script should be referenced with a METHOD of POST.\n");
      printf("If you don't understand this, see this ");
      printf("<A HREF=\"http://www.ncsa.uiuc.edu/
           SDG/Software/Mosaic/Docs/fill-out-forms/
           overview.html\">forms overview</A>.%c",10);
      exit(1);
   }
   if(strcmp(getenv("CONTENT_TYPE"),"application/x-www-form-urlencoded")) {
      printf("This script can only be used to decode form results. \n");
      exit(1);
   }
   cl = atoi(getenv("CONTENT_LENGTH"));
```

Figure 117 (Part 1 of 2). post_query Code

```
   for(x=0;cl && (!feof(stdin));x++) {
      m=x;
      entries[x].val = fmakeword(stdin,'&',&cl);
      plustospace(entries[x].val);
      unescape_url(entries[x].val);
      entries[x].name = makeword(entries[x].val,'=');
   }

   printf("<H1>Query Results</H1>");
   printf("You submitted the following name/value pairs:<p>%c",10);
   printf("<ul>%c",10);

   for(x=0; x <= m; x++)
      printf("<li> <code>%s = %s</code>%c",entries[x].name,
           entries[x].val,10);
   printf("</ul>%c",10);
}
```

Figure 117 (Part 2 of 2). post_query Code

You can note that there is two comparisons: one with the REQUEST_METHOD environmental variable to see if this is working with the POST method or something else, and the second to see the content type. The third variable we check is the CONTENT_LENGTH. We convert its contents from an ASCII string to an integer and put the answer into the cl variable.

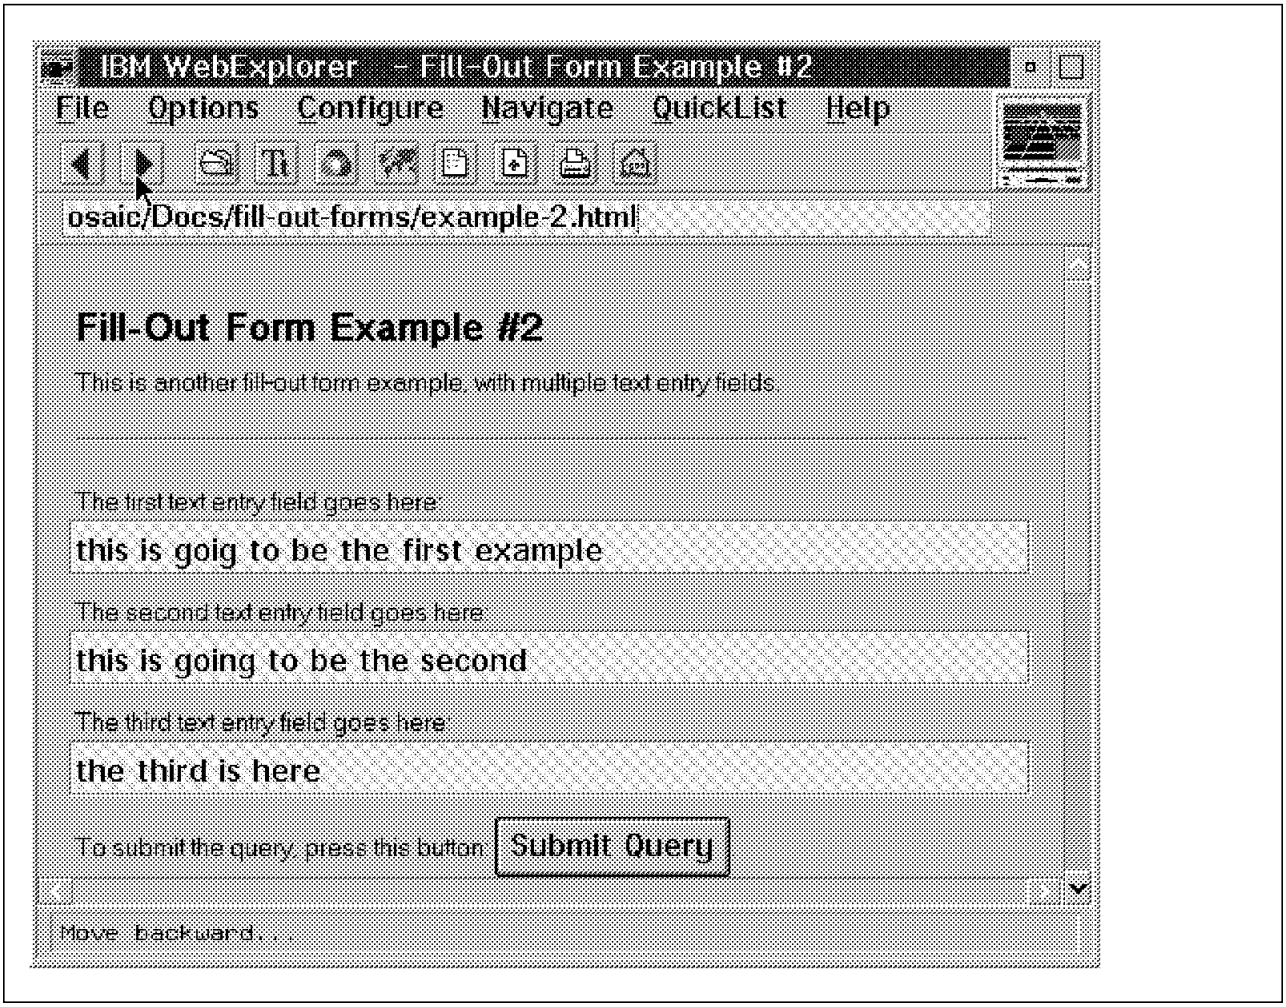


Figure 118. Input from post-query.c. This is the same output as in query.c.

For decoding this part we use fmakeword, indicating the standard input as the main file and cl, which is going to be modified to get the new length after the function call, to give the number of bytes to use.

We use the plustospace, unescape_url and makeword to finish the decoding like we did in the query.c example.

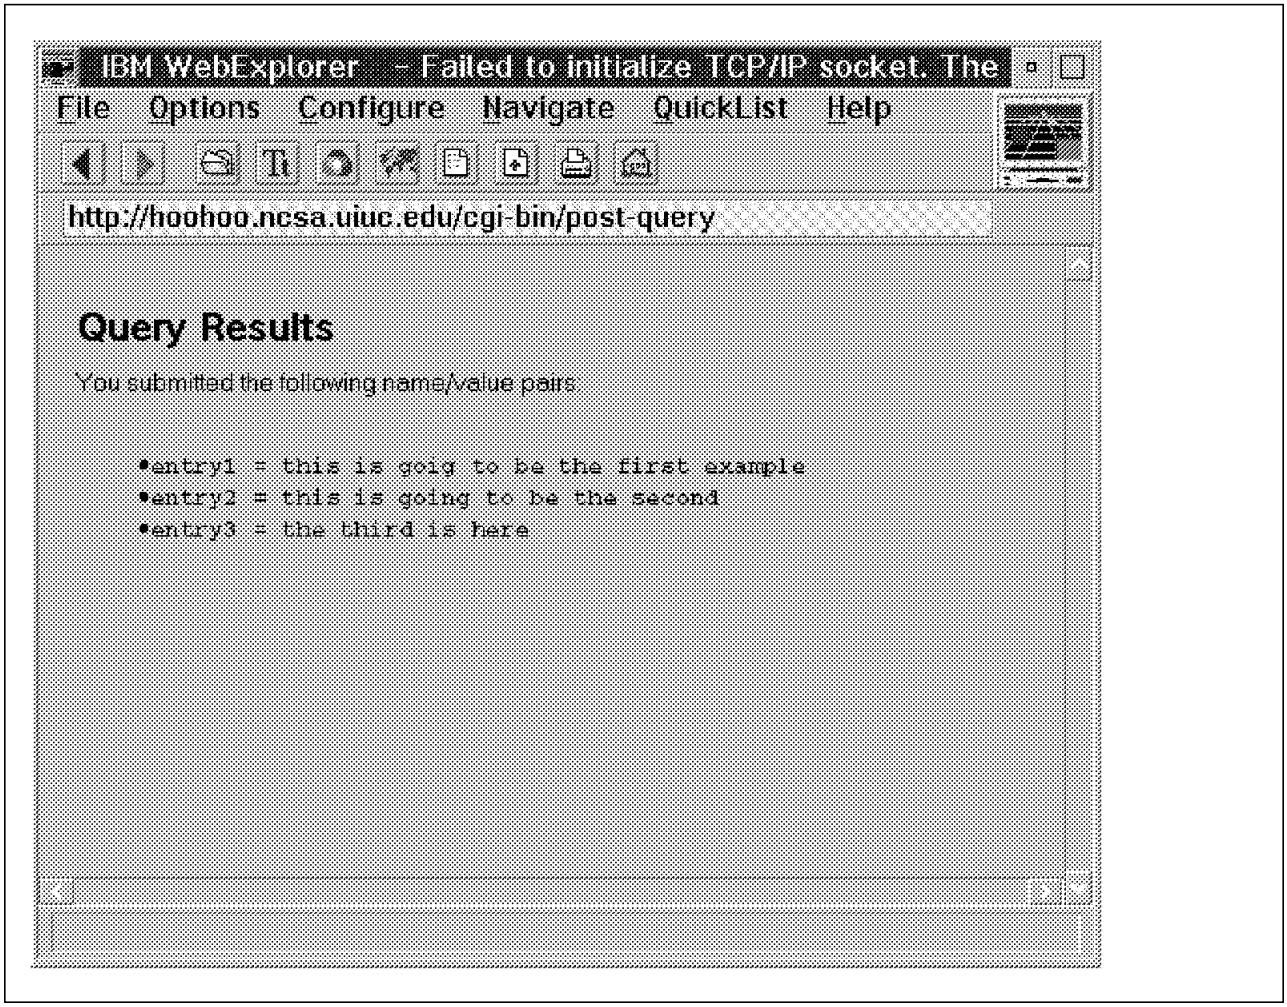


Figure 119. Output from post-query.c.. This is the same output as in query.c.

4.5.10 Ideas for Interesting Pages with CGI Programming

The CGI programming is one of the more powerful tools for the intranet and the Internet. Examples of what is posible with the CGI include:

- **The internet virtual yellow pages chat** http://www.vyp.com/cgi-bin/chat/login
- **Yahoo search engine** http://www.yahoo.com
- **IBM′s Infomarket search engine.** http://www.infomkt.ibm.com/pubbin/imsQuery?immfmt3=ht3
- **Virtual pizza ordering.** http://www2.ecst.csuchico.edu/~ pizza
- **A good place to send electronic postcards for free.** http://postcards.www.media.mit.edu/Postcards

4.5.11 Error Handling with CGIs

Another useful solution for a CGI program is the error handling, In the CD-ROM you will find the list of a CGIs that make this possible. This file is also from NCSA and it's PERL written. For the error handling you will note the use of other environmental variables.

4.5.11.1 REDIRECT_REQUEST

This is the request the client did exactly like the server received it.

4.5.11.2 REDIRECT_URL

This is the URL that caused the error (if a CGI is not working properly it can return an error, so the REDIRECT_REQUEST and REDIRECT_URL are not going to match).

4.5.11.3 REDIRECT_STATUS

This is the default message the server should return.

Some servers use the standard NCSA configuring file convention and use the srm.conf file in order to configure the URLs that have to be called in case of error. For further documentation you can visit:
`http://hoohoo.ncsa.uiuc.edu/setup/srm/Overview.html`

4.5.12 CGI Security Considerations

The security of the Web server is a big concern when the company data is playing a roll on the application. The information and the server are an important part of the company.

The first thing you have to observe is to make the NOBODY user run your server if you are on a UNIX like system (is part of the configuration tools on the IBM server).

Second, if you make variables to use char strings you have to be careful the way you make your program, and a UNIX-like system if the server is not secure enough the nobody user can overpass the limit of your strings and open a remote shell by itself. Most of the new servers don't have this security flaw anymore. In order to prevent this you have to use dynamic variables or use huge static variables (remember we said to use char strings with a 32 KB size or more).

Don't allow the users (unless you real need to) to make command line operations (which can be done using the system(), popen(), or REXX interpret) do anything, This is a big security hole in your CGI program (not from the server).

4.6 Virtual Reality Modeling Language

3-D graphics have become one of the most researched areas in computer science in the last few years; techniques such as radiosity and ray tracing are popular among computer science engineers.

The Open Graphics Library, also known as Open GL has opened a good way to create cross-platform programs to take advantage of the operating system and hardware capabilities and perform better graphics with the same interface.

Systems such as OS/2 Warp Merlin and Windows NT (Version 3.5 or later) have native support for this library. Open GL has been a good tool to perform 3D graphics and the internet has taken advantage of it.

VRML is a language that allows the programmer to create 3-D objects, link the files and create a common browser language to navigate in 3-D worlds. VRML needs (as HTML) a special browser to display these virtual places. Open GL allows you to create browsers in better shape, but it is not the only resource used by programmers, most of the companies use their own engines. Companies such as Eagen use Open GL to create their browsers. Eagen has developed warpspace, a VRML browser that will work with IBM Web Explorer. Other browsers for the rest of the platforms are available, and the Netscape Navigator has bundled a VRML engine on it.

The way warpspace works is loading only those files that use the VRML 1.0 specifications, parsing them and using the Open GL engine for OS/2.

Netscape bundles superscape that is launched when a VRML file is found. Netscape has VRML across different platforms such as Windows 3.1, Windows 95, Solaris, HP-UX and AIX (a version of OS/2 has been announced for the last quarter of 96, just when Merlin arrives to the market).

4.6.1 VRML specifications

VRML is a language that is based on solid construction graphics and uses a syntax based on nodes that represents objects in a virtual world and the actions that you can do on it. For a node an action can represent an anchor to another world or other HTML page, which where the information based on HTML pages and the virtual worlds can meet.

To create VRML worlds you must have a basic knowledge of Computer Graphics and solid construction images, if you don′t want to be assisted by an authoring tool. Authoring tools are very simple to use and can help to create a good impression, but the more knowledge you have about the language and computer graphics the more you will be able to create on your virtual world page.

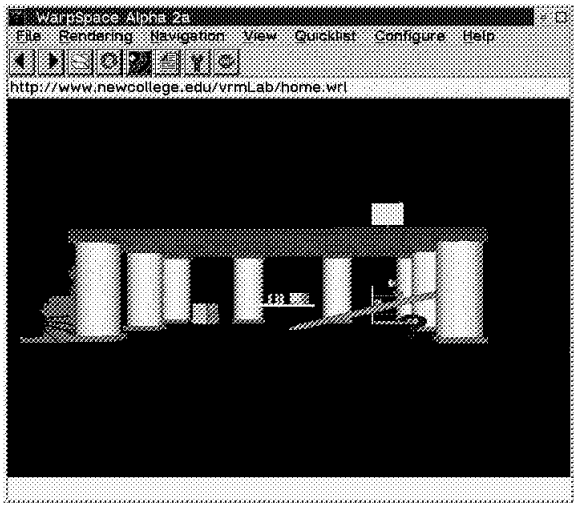


Figure 120. WarpSpace is a VRML browser for OS/2. It works fine with WebExplorer with a very simple configuration.

If you do, you will understand the following VRML basic program and you are ready to learn the language.

```
# Just a VRML example
Separator {
        Transform  { translation 20 -100  80}
        Material {
                 difuseColor 0 0 1
                 }
        Sphere {
                 radius 17
                }
}
```

If you do not understand the previous code above we give you a brief explanation of it. VRML works on 3-D coordinates and 3-D objects, where every separator is an object that is painted on the browser. You have to tell the browser the material for every object and the correct transformations (the transformations are translations, rotations and scales) in the correct order. It is not the same result if you make a translation and a rotation than if you do it backwards. After that, you specify the object to paint (for example, if you are painting an sphere with a 17-unit radius with the center at 20,-100,80 and a material that has a blue color).

To get more information about the language you can visit:

`http://webspace.sgi.com`

In this place you can find complete information about the VRML 2.0 language, the complete syntax, and a complete learning manual with about 90 pages of explanations and examples. This is the right place to look if you are thinking of learning VRML. This place has been developed by the people who developed VRML: Silicon Graphics.

To get information about Warpsace go to:

`http://www.eagen.com`

To get information about Netscape and Superscape go to:

http://www.netscape.com

In Yahoo (http://www.yahoo.com) you can find some authoring tools to create VRML pages in a simpler way to create your own virtual world with the limitations on the tool you use.

In case you are planning to get a VRML world we recommend that you to create alternative pages for those people who don't have a VRML browser.

One tool that can help you create this virtual world is Virtual House Builder. It runs under Windows 3.x or Win/OS2 and gives you a view of the objects while you are constructing your documents. It is easy to use and easy to learn.

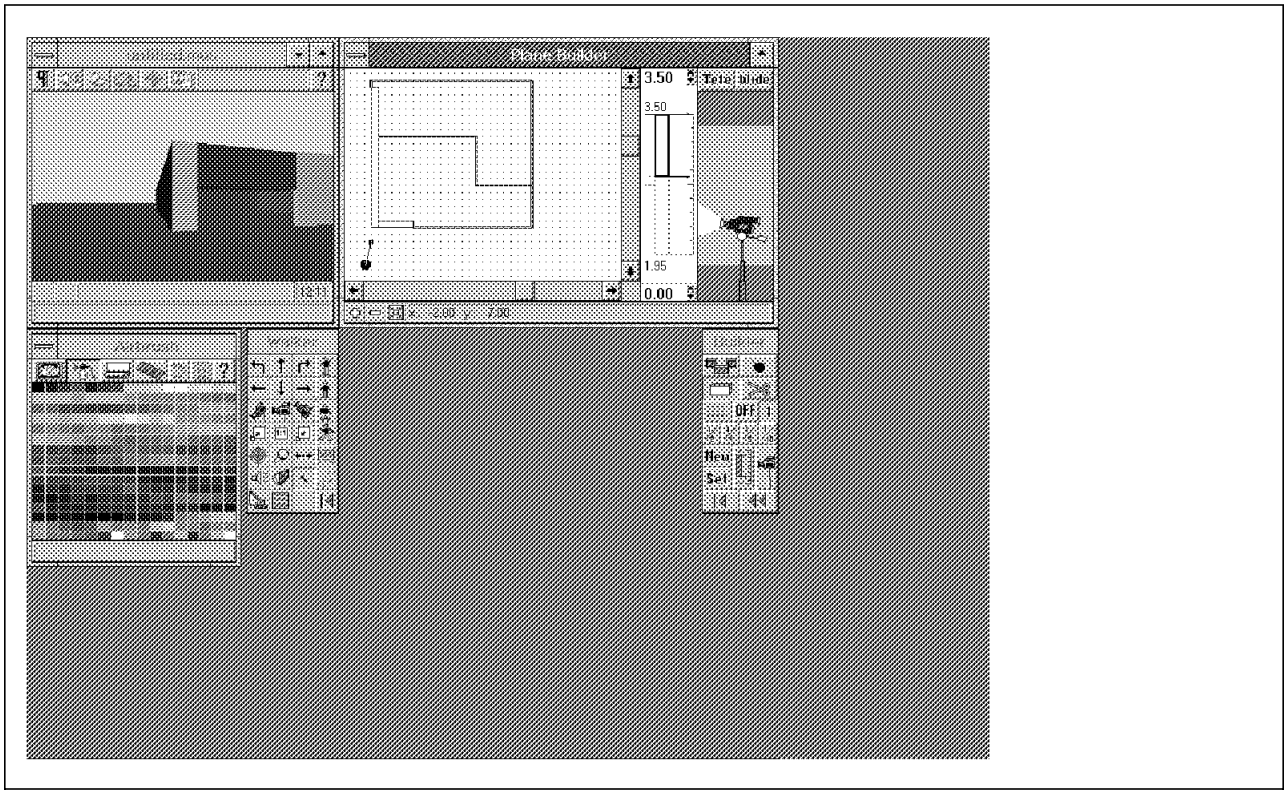


Figure 121. Virtual House Builder Helps You to Create VRML Worlds.

The way Virtual House works is very intuitive and allows you to create and walk around objects, most of them boxes. For more information you can go to:

http://www.paragraph.com/

4.7 Other Useful Tools

One of the improvements that the software producers have done to the CGI interface is to add APIs (Application Program Interfaces). Most of them are special functions that allow the programmer-made CGIs in an easier way. Apache, IBM Server, Netscape Commerce Server and others have their own set of APIs that allow you to have the same server to listen to requests from different IP addresses, create specific user directories and attend some specific requests, give a better authentication and provides good security.

You can find more information about these servers at:

- Apache:

 http://www.apache.org

- Netscape Commerce Server:

 http://www.netscape.com

- IBM Internet Connection Server:

 http://www.ics.raleigh.ibm.com

4.7.1 JavaScript

JavaScript is a language introduced by Netscape to use as a part of the Netscape Navigator extensions. This language is different from the Java language and it is more restricted in the things you can make with it. JavaScript pretends to be a useful tool to create interactive pages and in-line calculations, helping the submission forms to perform a better interface.

The Java Script functions have to be written in the HTML file as a part of the HEAD section. After the <head> tag you most use the <SCRIPT LANGUAGE="JavaScript"> tag. You must begin to write your code here, but be sure this code is going to display in other non-Netscape browsers (Oracle is also supporting JavaScript); to prevent this you must put all your code in HTML comment the first line of comment will be ignored like a part of the script.

JavaScript allows you to interact directly from the user; the language was developed to create interactive pages and enhance the interaction between Netscape's plug-ins and the Java Applets.

A good place to start to learn the language is:
http://home.netscape.com/comprod/products/navigator/version_2.0/script/script_info/index.html

Other similar languages were developed with the same idea, such as the Visual Basic Script, that runs over Microsoft Explorer. Microsoft wants to make the VisualBasic a standard on the Internet world and is making the applications compatible with this new philosophy. Visual Basic Script is similar to the Visual Basic language.

JavaScript is useful for applications that require interaction between the user and the server, and the server wants to make a specific process for itself. If you have something to update in a form and you want to do it in real-time, you must use JavaScript; if you have an application that does not require additional interactions with the server, you can use Java-like applications with animations.

The best way to find out what can be your best choice in the language selection for your application is to learn both languages and go with your feelings. Some examples of applications with JavaScript are in the address above.

Chapter 5. Java Programming

Java is important because it brings to the computer society the binary compatibility that has been requested for a long time.

All operating systems are incompatible with each other,including even programs for the same operating system on different hardware platforms.

Sometimes this can be fixed with a standard language supported on all platforms (such as C and C++). You only have to use ANSI C code to make it portable, so you couldn't make anything with the GUIs. The problem with interpreted languages was even worst, having no standardization (REXX has already an ANSI standard) and no GUI code portability.

Java creates the concept of *bytecodes*, which is a similar concept to the Virtual Machine on VM or the DOS Virtual Machine on OS/2. This translates from a set of codes previously declared (the API from DOS or the VM API) to the proper code for the operating system. Java has a Java Virtual Machine running in the operating system that responds to a code that is very similar to those on the computer processors That's why you have to compile it, and after that it has to be interpreted. The interpreter makes the translation faster than the regular interpreters because the classes (applications or applets) are in a similar code as the machine's.

The improvement on this is very simple. Now you have something very similar to a binary compatibility. Your code runs the same in OS/2, AIX or the Windows 32-bit family without recompiling it or changing something in the GUI code to keep the look and feel in all platforms.

Java also provides a natural way to make object oriented programming and one interface specially created to make applications for the World Wide Web using the browsers and extending the HTML language with the <Applet> tag.

5.1 Applets and Applications

Java is more than a tool to create *cute* pages on the WWW. It can be a tool to make client/server applications and stand-alone applications as well.

The applications that already have the ability to run in a browser are called applets.

The applications are not restricted in any way. You can do anything you want. You can run programs that read and write files, can make communications between two different machines (or more) using any port (using TCP/IP) and program your own protocol.

When you are writing applets you are working in a restricted place.

© Copyright IBM Corp. 1996

5.1.1 Applets Security Restrictions

Sun allows people to try to break the security on both sides (server and client) of the applets in order to improve it. The restrictions are:

1. Applets can not read or write from the file system. Except for those directories that the user defines in an access control list, it is empty by default. This list is specific for the browser you use, some browsers will not be allowed to read or write on the file system at all.
2. Applets can only communicate with the server where the applet was stored. This restriction can also be avoided by the browser, so you can't count on it.
3. Applets can not run any program on the client system. For all UNIX systems this also includes forking a process.
4. Applets can not load DLLs or native programs to the local platform.

As you can see, almost all the security that Java provides is client-focused, so if you are planning to make an applet, you have to see about your server security. This is very important if you are planning to establish a communication between the client and the host. Avoid this approach if it is possible.

5.2 Java Basics

If you are not familiar with the class, object oriented language or any other term we use in this chapter, visit the following URL: `http://java.sun.com/tutorial`.

To obtain the latest versions of the Java Compiler or the Java Development Kit drop by:

- `http://java.sun.com` for Sun solaris, Win 95 and NT and MacOS.
- `http://ncc.hursley.ibm.com/javainfo` for the AIX and OS/2 versions. There will be OS/390 and OS/400 versions soon.
- `http://hpcc998.external.hp.com/gsyinternet/technology/java/JDK.html` for the HP/UX version.

Java is a full object oriented language. The minimum compilation unit is a file; one file should have at least one class.

The Java language provides structural programming interfaces to create the methods. It seems a lot like C++.

If you already are a C++ programmer, you are going to have one advantage, but try to forget everything about C++ outside the object-oriented approach such as global variables, functions and procedures (void functions). Don't try to convert your Java language programs into a C++ extension. They have differences, and those differences are big enough to prevent you from trying to write Java programs thinking like a C programmer.

5.3 Differences between Java and C/C++

Java inheritance is simple; in C++ we have multiple inheritance.

There are no pointers on Java, but you can reference any previously declared object, so you have all the power of pointers without all the confusing C/C++ syntax for pointers.

If you need to call a specific OS function, you can load it, but remember: if you load specific OS functions, DLLs or programs, you are not making independent platform applications and you depend on the OS or the program in order to get good program behavior.

The arrays in Java are a special class of objects; this means that they are not *memory address* references like they are in C, and they have their own methods. The multidimensional arrays in Java are, as a matter of fact, arrays containing other arrays.

Strings are also a first class object; they are not null characters terminated like they are in C, so you don't have to worry about the array size or getting out of bounds in your string code.

Java has a garbage collector built in. This means that you don't have to make explicit disposals of the objects; Java interpreter will do it for you. If you really want to do it then make your object equal to null.

No struct, union or typedef keywords are used. Remember, you are working with objects; you are not working with structures anymore.

The Java language provides *platform independence* on data types. This means that a char is going to be a 16-bit data type; the int is a 32-bit data type in all platforms.

The only unsigned data type is the char. Java does not allow operator overloading like C++. We discuss other differences between Java and C++ later.

5.4 Java Compiler and Interpret

Before starting with the language itself, let's see how you should compile programs and classes.

The compiler name is *javac* and you use it from the command line followed by the name of your file.

In order to compile, check out this list of things to remember:

1. The file and the main class should have the same name.
2. Some operating systems are name case-sensitive.
3. Java is case-sensitive.
4. The extension for the file should be .java.
5. The extension of the compiled programs are .class.
6. To run a compiled class use java followed by the class name. Remember Java is case-sensitive.
7. To run an applet, make the HTML file and open it from your browser.

You can read the word compiled class, but we are talking about an interpreted language. The fact is, we call it a compiled class because you are generating final instructions, like any compiler does. The difference is that you are running this final code in a Java Virtual Machine. To put it another way, you are emulating a special class of processor and system that does not exist.

5.5 Language Syntax

In the Java language you will have a declaring section where you will declare all the objects you want to use in your programs and your classes code segment. We get back to this later.

Your classes have to have their properties declaration segment and the methods code segment.

Look at the traditional Hello World Application in Figure 122.

```
class hello
{
  public static void main(String args[])
  {
    System.out.println("Hello world!");
  }
}
```

Figure 122. Java Hello World Application. The typical Hello World application must be saved in a a hello.java file.

As long as you don't need other classes or variables, your only declarations are your object class and the main method.

5.6 Variable Types and Declarations on Types

The Hello World application could also be written as shown in Figure 123.

```
class hello
{
  public static void main(String args[])
  {
    String myString = "Hello World";
    System.out.println(myString);
  }
}
```

Figure 123. Other Hello World. The typical Hello World Application must be saved in a hello.java file.

As you can see, you must do the variable declarations naming the type or object name before the variable name; you can also do the assignment at the same time.

Another important note is that you must use the semicolon (;) symbol if you are writing a sentence and you can create sentence blocks using the bracket ({}) symbols.

The primitives data types for Java are defined in Table 21.

Table 21 (Page 1 of 2). Java Data Type Specifications

Data type	Definition
byte	8 bits signed (-128 to 127)

Table 21 (Page 2 of 2). Java Data Type Specifications	
Data type	**Definition**
char	16 bits unsigned, using Unicode character set
short	16 bits signed (-32768 to 32767)
int	32 bits signed (-2,147,483,648 to 2,147,483,647)
long	64 bits (-92,233,720,036,854,755,808 to 92,233,720,036,854,755,807)
boolean	Only for `true` and `false` values
float	32 bits single precision IEEE 754 compliant
double	64 bits double precision IEEE 754 compliant

When you are going to use attributes, you may declare them just before the methods. For example:

```
public class sum{
public static void main(String args[]){
int a;
float b,c;
  a=5;
  b=6.0f;
  c=(float)a;
  System.out.println(b+c);
}
}
```

Figure 124. sum.java File. sum.java makes a sum of two numbers and casts some variables.

In order to assign the values between variables or literals (which are constant values), you have to cast the variables.

When you are creating methods that are not the main methods, you can refer to the class attributes using the this command followed by the dot operand and the name of the attribute. In order to finish this part, you have Table 22 and Table 23 on page 241.

Table 22 (Page 1 of 2). The Operators List	
Operator	**Meaning**
Arithmetic operators	
+	Addition
-	Subtraction
*	Multiplication
/	Division
%	Modulus
Assignment Operators	
=	Assignment
x + +	Return the x value and increment it
+ + x	Increment the value of x and return the value
x--	Return the x value and decrement it

Table 22 (Page 2 of 2). The Operators List

Operator	Meaning
--x	Decrement the value of x and return the value
+ =	Add and assign
- =	Subtract and assign
/ =	Divide and assign
* =	Multiply and assign
% =	Modulus and assign
& =	And and assign
\| =	Or and assign
> > =	Right shift and assign
> > > =	Zero fill right shift and assign
< < =	Left shift and assign
^ =	XOR and assign
Bitwise operators	
\|	OR
&	AND
^	XOR
< <	Left shift
> >	Right shift
> > >	Zero fill right shift
~	Complement
Comparison operators	
>	Greater than
<	Less than
> =	Greater than or equal to
< =	Less than or equal to
= =	Equal
! =	Not equal
&&	Logical AND
!	Logical NOT
\|\|	Logical OR

If you are a C programmer, you already know that the variable type declaration and the operators are the same.

In Java you cannot overload the operators. This means that your objects cannot have an operator for them and you cannot make an addition to them. In order to create something like this, you will have to create your own methods instead of operators.

To make a reference to an attribute or a specific method of the object, you use the dot (.) operator.

Table 23. Operators Precedence. The first one has the highest priority.
Operator
. [] ()
+ + -- ! ~ instanceof
new (castingtype) expression
* / %
+ -
< < > > > > >
< >
= = ! =
&
^
\|
&&
\|\|
= += -= *= /= %= ^= &= \|= < < = > > = > > > =

The String object is one special class and is the only one that has an operator by itself. You can concatenate strings using the addition (+) operator, and you can concatenate all kinds of objects with the same operator if you have immplemented the toString method. This method allows an object to be converted to a String object.

To create a better view of this, look at the example shown in Figure 125 on page 242.

```
class person{
String name;
String Lastname;

public String toString(){
       return "You are talking about "+name+" "+Lastname;
     }

void Setname(String name){
  this.name=new String(name);
}

void SetLastname (String Lastname){
  this.Lastname=new String(Lastname);
}

static public void main(String args[])
{
 person Me=new person();
 person Nilson=new person();
 Me.Setname("Carlos");
 Me.SetLastname("de Luna");
 Nilson.Setname("Nilson");
 Nilson.SetLastname("Baptista");
 System.out.println(Me);
 System.out.println(Nilson);
}
}
```

Figure 125. A Simple Class. The person class is going to print its name.

You can see Figure 125 has things such as the creation of an object and the call to their methods.

5.7 Classes, Objects, Inheritance

Java classes are the representation of the objects you want to create. Those objects have their own behavior depending on the methods you use and the information (parameters) you use with them. When an object can be loaded directly from the compiler, such as an applet or a program, it's because the class has implemented some special and useful methods to run the application or applet.

For an application, as you can guess from the examples above, the main method is the procedure that the compiler is going to find or execute. The applet is something we discuss later in this chapter.

The Java objects have simple inheritance, which means they can have inheritance only from one class, but it doesn't mean you can't use more than one class. To use the other class you have to declare them before your class declarations. Sometimes you need a different method act if you call it with one or two parameters. The only thing you have to do is write those methods with the correct number of parameters on them. This is called overloading a method.

When you create a new object (as you can see in Figure 125) with the new operator, you call it a constructor method. This is a method that is going to be

called when an object is created. The name of that method should be the same as the class name. In Figure 126 on page 243 you can see an example of this and how you can override the toString method.

```
class person2{
String name;
String Lastname;

public person2(String name){
this.name= new String(name);
}

public person2(String name, String Lastname){
 this.name= new String(name);
 this.Lastname = new String (Lastname);
}

public String toString(){
       return "You are talking about "+name+" "+Lastname;
     }

void Setname(String name){
  this.name=new String(name);
  this.Lastname=" ";
}

void SetLastname (String Lastname){
  this.Lastname=new String(Lastname);
}

static public void main(String args[])
{
 person2 Me=new person2("Carlos");
 person2 Nilson=new person2("Nilson", "Batista");
 System.out.println(Me);
 System.out.println(Nilson);
}
}
```

Figure 126. The person2 Class. The person2 class has 2 constructors and overrides the toString method.

Sometimes you will need the class to do something before the object is eliminated. In order to tell the object what to do before it dies you must write a destructor method. The destructor method in Java is called finalize().

5.8 Overriding Methods

When you inherit from a class, you are getting all the methods that exist on it. But if you need a different behavior from one of those methods you should again write the method that the new class has to follow.

The examples in Figure 125 on page 242 and Figure 126 show you how to override the toString method. All classes have this method by default, which returns the name of the class followed by empty brackets.

As you can see, the toString method should return a String object; you use the return keyword to do this. But if you need to make reference to a method that

was overridden (that means a superclass method) you must use the dot (.) operator to indicate the class which you are making the reference to and the name of the method with their parameters.

5.9 From Arrays to Loops

The arrays in Java are special kinds of objects (the same as the strings). You can make an array containing any class of type or object. As long as the arrays are objects you cannot make reference to them the same way C or C++ does; you have to make the reference to the value that you are looking at. An array does not represent a memory allocation place.

In order to declare an array you must put the type of the array followed by the brackets containing the length of it. If you want to make a dynamic array, you must leave the brackets empty. You can also declare them using the type or object name followed by the empty or not empty brackets and the name of the variables. For example:

```
int myIntegerArray[20];
int[] myDynamicIntegerArray;
.
.
myDynamicIntegerArray= new int[theDynamicLength];
```

You can also assign their value using the block separator chars:

```
int [] pairs={2,4,6,8,10,12};
```

As long as Java does not support multidimensional arrays, you create arrays containing arrays, emulating the dimensions, using the same syntax of C:

```
String cellContents[10][10]
```

5.9.1 Casting Elements

When you are writing a program one of the problems you have is using different types of variables (sometimes a float, sometimes an integer) and trying to use those kind of variables together. In order to do this you can cast the variables to fix the correct type. To do this, you indicate the type you want to have in parenthesis and the variable name you want to convert outside the parenthesis. For example:

```
int a;
float b;
.
.
a=(int)b;
.
.
```

You can also make a casting of objects from a subclass to a superclass.

5.9.2 Conditionals

In order to create changes in the flow of your method execution procedure, you must use conditionals or loops. In this part you can see the if...then...else... conditional; after that you will see the switch... case... default... and the loops.

The if conditional is used to execute a part of a method only if the condition between the parenthesis has a true value. In Java the true value is one of the two possible states for a boolean variable; you cannot have an if that answers to a numerical value (like you do in C and many other languages).

For example:

```
.
.
if (rainChances>50) System.out.println("Today is going to be a wet day");
.
.
```

If the condition is true, the statement for printing is executed. You can have a block of instructions instead of one:

```
.
.
if (rainChances>50)
  {
   wetDays=wetDays++;
   System.out.println("have been "+wetDays+" wet days on the year");
  }
.
.
```

Sometimes you need to do different things if the condition is true or if the condition is false. To do this use:

```
.
.
else clause for this cases:
if (rainChances>50)
        System.out.println ("Another wet day is waitting outside");
else
        System.out.println ("It will be a day without rain for me...");
.
.
```

5.9.3 switch... case... default

Use the switch command when you have a multiple condition strategy with a simple variable. Put the variable in parenthesis after the command and use a block to order your strategy.

Inside the block put all your possible cases using the case keyword, followed by the statements in the proper case. Use the default keyword for those cases that are not considered with the case keyword.

For example:

```
.
.
switch (day) {
      case 1:
             System.out.println("Monday");
             break;
      case 2:
             System.out.println("Tuesday");
             break;
      case 3:
             System.out.println("Wednsday");
             break;
      case 4:
             System.out.println("Thursday");
             break;
      case 5:
             System.out.println("Friday");
             break;
       case 6:
             System.out.println("Saturday");
             break;
       case 7:
             System.out.println("Sunday");
             break;
      default
             System.out.println("Invalid day on Gregorian calendar");
             }
.
.
```

Look at the break words after each action is finished. Use it to prevent the code from following the instructions (corresponding to the next condition). If you want to make the same code for similar conditions you can write it only once:

```
.
.
switch (month){
      case 1:
      case 2:
      case 3:
           QuarterProfit[1] = QuarterProfit[1]+MonthProfit[month]
           break;
      case 4:
      case 5:
      case 6:
           QuarterProfit[2] = QuarterProfit[2]+MonthProfit[month]
             break;
      case 7:
      case 8:
      case 9:
           QuarterProfit[3] = QuarterProfit[3]+MonthProfit[month]
             break;
      case 10:
      case 11:
      case 12:
           QuarterProfit[4] = QuarterProfit[4]+MonthProfit[month]
             break;
      default
             System.out.println("Invalid month on Gregorian calendar");
```

```
            }
.
.
```

In the second example add the monthly profit to the corresponding quarter profit generating less work than in languages such as BASIC where you have to write it each time for every month.

5.9.4 do...while

This is the statement you use to obtain loops where the inside code has to be done at least once. The syntax is:

```
do {
       looping statements
   } while(continuing condition);
```

The loop will be executed while the condition inside the while parenthesis has a true value.

5.9.5 while and for Commands

The while and for loops can or cannot be executed depending on the value of their conditions; both of them work the same. (If you are a C programmer you know how they work.)

In the while loop you have to follow this syntax:

```
while (continuing condition) {
      looping statements
      }
```

where the continuing condition has to be true to execute the looping statements. These statements are executed until the condition stops being true.

The for statement has the following syntax:

```
for (initial expression; continuing condition; looping operation) {
       Java statements
       }
```

At the beginning of looping the for statement, execute the initial expression. This is used for an initialization rather than a common expression. Before doing the looping the looping condition is evaluated. If it's true, then the loop is executed. Every time the loop is completed, the looping operation, normally used to increment or decrement a variable, is executed and the condition is evaluated. (Remember that these conditions are boolean variables, not integer variables such as in C.) In order to make these examples clearer, you have three examples shown in the Figure 127 on page 248, Figure 128 on page 248 and Figure 129 on page 248 that make a counting output from 1 to 10.

```
class doclass{
public static void main(String args[]){
    int i;
    i=1;
    do {
        System.out.println(i);
         i++;
         }while(i<11);
 }
}
```

Figure 127. An Example With do. doclass shows the use of do to create a 1 to 10 counting.

```
class whileclass{
public static void main(String args[]){
    int i;
    i=1;
    while(i<11) {
        System.out.println(i);
         i++;
         }
 }
}
```

Figure 128. An Example With while. whileclass shows the use of while to create a 1 to 10 counting.

```
class forclass{
public static void main(String args[]){
    int i;
    for(i=1;i<11;i++)
        System.out.println(i);
 }
}
```

Figure 129. An Example With for. forclass shows the use of for to create a 1 to 10 counting.

At the moment you run these three programs, you see the outputs are exactly the same. Do a countdown for exercise.

5.9.6 Labeled Loops

Almost all programming languages have implemented the goto keyword. Java does have this keyword like a reserved word of the language, but it does not have this function implemented.

In some languages, such as C, you put some labels and put the goto expression to send the program to that specific point. Java does not have that functionality.

In order to get out from a loop without executing the next statement you can use the break keyword (like you did in the switch command), but if you have a loop inside another loop, you may be want to get out of the outer loop rather than the inner one. To do this you must use labeled loops.

Create a labeled loop following these steps:

1. Create a label. Put the label name followed by a colon (:): For example:

   ```
   labeledline:
   ```

2. Put the break statement followed by the label name where you need it.

5.10 Applets Basics

The applets are very special applications. They normally use the GUI to create a better interaction, but they have restrictions to keep the security on the Internet.

In the applications above we haven't used the GUI at all; we use it when we work with applets and also implement the inheritance.

An applet is a subclass of the panel class, which is a subclass of the container class.

The hierarchy map is shown in Figure 130 and shows you a little about the implementation of the Java language.

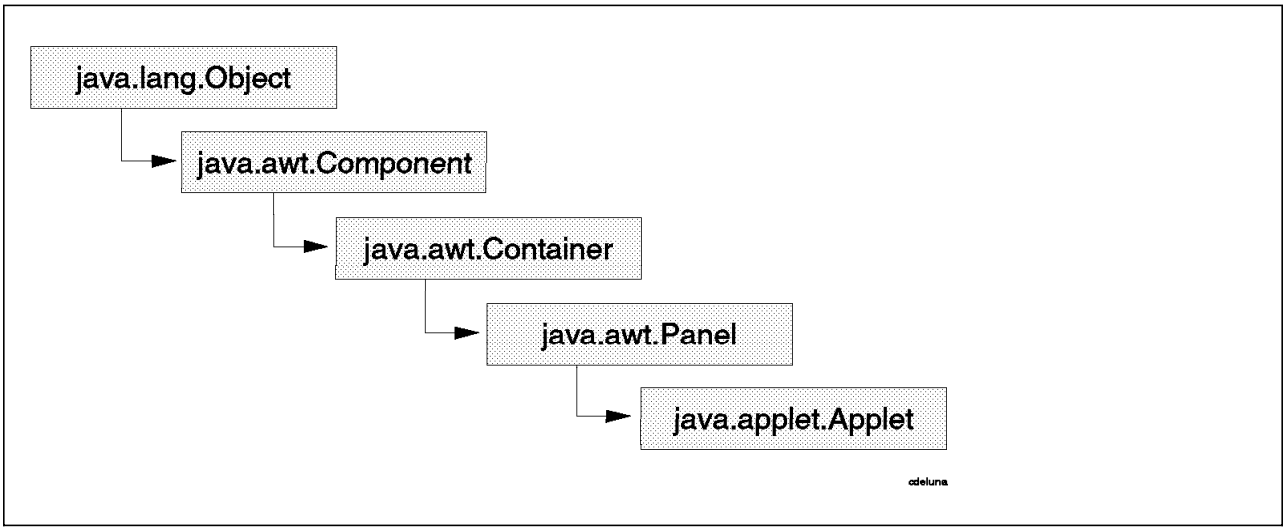


Figure 130. Java's Applet Inheritance Tree

To create an applet you create a subclass from the applet class and override some of the methods.

The thing you have to keep in mind is that all graphic libraries are stored as a part or a subclass of the AWT (Abstract Window Toolkit) class.

When you make an applet, the standard output is not the applet area presented by the browser or the applet viewer; it should be the Java console or the browser status line.

In order to put something in the applet area you have to draw it, making the font selection.

```
<HTML>
<Title>Clock Applet </Title>
<BODY bgcolor="#FFFFFF">
<APPLET Code="appletworld.class" WIDTH = 200 HEIGHT=30 ALIGN=RIGHT>
The clock is not displayed because you are not using an Java anabled
Browser.
</Applet>
</Body>
</HTML>
```

Figure 131. HTML File for the appletworld Class

```
  import java.awt.Graphics;
  import java.awt.Font;
  import java.awt.Color;

public class appletworld extends java.applet.Applet{

Font f = new Font("TimesRoman",Font.BOLD,12);

public void paint(Graphics g){
 setBackground(Color.white);
 g.setFont(f);
 g.drawString("Hello Applet's World!",5,20);
 }
}
```

Figure 132. appletworld Class Program



Figure 133. Screen with the appletviewer and Web Explorer with the applet.

As you can see, you have use an HTML page to see the applet. Using the <Applet> tag, program the Java class and run the sample loading your HTML file from a browser or with the applet viewer included with the Java Development Kit. To run the applet viewer use:

`appletviewer htmlfile` for UNIX systems

`applethtml file` for OS/2

You can see that there is some text outside the <Applet> and </Applet> tags. This text is recognized only on those browsers that are not Java-enabled.

In this example we are only overriding the paint method. This is called when the applet apears or when the repaint() function is called.

You are creating the subclass from applet when you put the extends keyword in the program. Another interesting point is that the class has to be public. If the class is not public you cannot use it to create an applet.

5.11 Implementing a Simple Clock

The next example shows one easy applet, a clock. With the first example you see how to use a simple program and after that you see how to use threads.

```
  import java.awt.Graphics;
  import java.awt.Font;
  import java.awt.Color;
  import java.util.Date;


public class nothreadclock extends java.applet.Applet{

Font f = new Font("TimesRoman",Font.BOLD,12);
Date d1;

public void start(){
while(true){
        d1 = new Date();
        repaint();
       }
}


 public void paint(Graphics g){
  setBackground(Color.white);
  g.setFont(f);
  g.setColor(Color.blue);
  g.drawString(d1.toString(),5,20);
 }
}
```

Figure 134. The Clock in an Applet. nothreadclock applet seems to be OK, but it is not.

The program shown in Figure 134 draws the hour in the same place over and over, so you will have the right hour displayed; but if you remain repainting

without sharing resourses from the system you will have nothing, and worst than that, you can crash the operating system. But the important part on the example is to show you that an important element on an applet is the `start` method.

In the applets you don't have to write the main method, unless you want the applet to be an application too. The reason is simple: you are not the main program, you are a part of a more complex program. The methods you have to write depend on what you want to do; the most common methods that you override are explained in the following sections.

5.11.1 The init Method

The init method is called when the applet is loaded or reloaded. The initialization part is here. If you need to load an image, create objects or set up your applications (depending on the applet parameters) you can do it here.

The init method should be a public method, and it does not return anything. Therefore the method should be overriden like this:

```
public void init()
```

The applets parameters are posted in the HTML file that makes the class reference. To put a parameter to your applet use the `<PARAM NAME=parametername VALUE=value>` tag.

To retrieve these parameters, use the getParameter method. The function is called with the name of the parameter and will return a string with the value. For example:

```
String parameter1=getParameter("Parameter1");
```

5.11.2 The Start and Stop Methods

After the initialization, the applet is started. This method can also be called when the page was unloaded and loaded again (that is, the moment when the applet starts). When the page is unloaded the applet stops, but it can be stopped by the programmer, in order to suspend the execution.

Both methods are public void type and they do not receive any parameters.

5.11.3 The Destroy Method

You can override the destructor method implementing public void destroy() but this is required only on special occasions.

This method applies only to the applets; to create a destructor method on any other object you need to override the finalize() method.

5.11.4 The Paint Method

You have to override this method in order to show something on your applet area. Here you draw all your stuff and load the images you need.

To override this method type `public void paint (Graphics g)` and remember to include the graphics class using the line: `import java.awt.Graphics` or if you want to use all the awt class you can use `import java.awt.*` to do it.

5.12 Threading Applets

Getting back to the example shown in Figure 134 on page 251 we have a simple way to fix it. The only thing you need to know, or at least understand, is what a thread is and how it works.

Even if we pause the main procedure, we will have all the control in one part and it will not work. When you use threads, you create small pieces of code running and sharing the resources. They are easier to control.

In order to use threads you must enable your class to run and create the code for every thread (shown in Figure 135 on page 254).

When you enable an applet to run, you program the public void run() method to be your main method.

```
  import java.awt.Graphics;
  import java.awt.Font;
  import java.awt.Color;
  import java.util.Date;


public class clock extends java.applet.Applet implements Runnable{

Font f = new Font("TimesRoman",Font.BOLD,12);
Date d1;
Thread running;

public void start(){
  if (running==null) {
    running=new Thread(this);
    running.start();
   }
 }

 public void stop(){
   if (running!=null){
      running.stop();
      running=null;
    }
  }

public void run(){
 while(true){
  d1=new Date();
  repaint();
  try{ Thread.sleep(1000); }
  catch (InterruptedException e){ }
 }
}

 public void paint(Graphics g){
  setBackground(Color.white);
  g.setFont(f);
  g.setColor(Color.blue);
  g.drawString(d1.toString(),5,20);
 }
}
```

Figure 135. The Clock That Works

You can see some new words have been added to the class declaration. When we add implements Runnable we are making this class available to use threads. We are incluiding a variable holding our applet (this).

The following are several methods:

- The start method stands to create the threads when those are necessary.
- The stop method stands to stop every thread if this is necessary.
- The run method is the main method for a thread; all the actions of this thread that are executing are here.

We put a thread to sleep without any problem, but if the thread stops in the middle of the dream or the applet (or just the tread) is destroyed, an exception appears. The try statement will be executed. If something goes wrong an exception will occurred and the statements on the catch block will be executed.

Look at how the Date class was used to create a new object called d1.

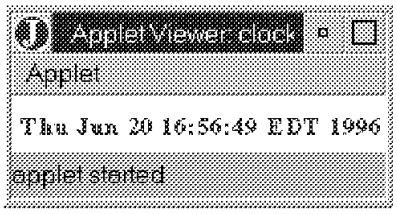


Figure 136. The Clock Applet in Action

5.13 Graphics on the Applets

The graphics are obtained using the Graphics class; this class allows you to draw bitmaps or GIF images.

The coordinate system is the common system in computer languages. (0,0) represents the superior left corner. There are only positive numbers representing the horizontal axis (in the first position of the coordinate) and the vertical axis.

Graphics class primitives are lines, ovals, rectangles, three dimensional rectangles, polygons and arcs. All figures can be filled or empty.

Text, using fonts, are also from the AWT, as you will see in Figure 137 on page 256 and Figure 139 on page 258

```
  import java.awt.Graphics;
  import java.awt.Font;
  import java.awt.Color;
  import java.util.Date;
  import java.awt.FontMetrics;

public class sign extends java.applet.Applet implements Runnable{
  Font f = new Font("TimesRoman",Font.BOLD,12);
  FontMetrics metrics = getFontMetrics(f);
  Date d1;
  Thread running;
  String text;
  int x;
public void init(){
   text=getParameter("text");
   if (text==null){
      this.text="Your HTML file is incomplete, the <Param> tag is mis    sing";
    }
}

public void start(){
   if (running==null) {
         running=new Thread(this);
         running.start();
    }
}
```

Figure 137 (Part 1 of 2). Flickering Sign

```
public void stop(){
    if (running!=null){
       running.stop();
       running=null;
       }
}

public void run(){
 while(true)
       {
         d1=new Date();
         repaint();
         try{ Thread.sleep(10); }
         catch (InterruptedException e){ }
        }
}

public void paint(Graphics g){
 setBackground(Color.white);
 g.setFont(f);
 g.setColor(Color.blue);
 g.drawString(d1.toString(),5,20);
 g.drawString(text,x,40);
 if ((metrics.stringWidth(text)+x)==0)
  {
    x=size().width /2;
  }
 else
  {
    x=x-1;
  }
}

}
```

Figure 137 (Part 2 of 2). Flickering Sign

The FontMetrics class is useful for getting information about the proportions of the font. The getFontMetrics help us to obtain that information on a specific font.

```
<HTML>
<Title>Sign Applet </Title>
<BODY bgcolor="#FFFFFF">
<APPLET Code="sign.class" WIDTH = 200 HEIGHT=60 ALIGN=RIGHT>
<PARAM Name=text Value="This is a typical test text for an applet like this">
The sign is not displayed because you are not using an Java enabled
Browser.
</Applet>
</Body>
</HTML>
```

Figure 138. HTML File for the Flickering Sign

The following methods are inherited by the applet and can be used for testing and setting colors on the working space:

```
setBackground(Color.white);
setForeground(Color.black);
```

The first one is used to change the background and the other for the foreground.

The color object has a constructor that allows you to get a specific color knowing the RGB codes.

```
color=new Color (R,G, B);
```

5.14 Animations, Sounds and Other Effects

The flickering produced in the sign class is due to the time that it takes the computer to paint every single pixel on the string and most of all in the screen, cleaning and painting on the screen. There are many solutions to the problem; the most simple is double buffering.

The double buffering is as simple as not erasing the screen, only repainting it.

To do double buffering you paint the shapes you need out of the screen When everything is finished you put it in the real world. In the Java class your applets have to override the update method. Without overriding update cleans the working area and repaints it. This method is called by the repaint method. In the following example you will see the correct sign class.

```
  import java.awt.*;
  import java.util.Date;

public class sign2 extends java.applet.Applet implements Runnable{
  Font f = new Font("TimesRoman",Font.BOLD,12);
  FontMetrics metrics = getFontMetrics(f);
  Date d1;
  Thread running;
  String text;
  int x;
  Dimension outDimension;
  Graphics outGraphic;
  Image outImage;

public void init(){
   text=getParameter("text");
   if (text==null){
      this.text="Your HTML file is incomplete, the <Parameter> tag is missing";
    }
}

public void start(){
   if (running==null) {
         running=new Thread(this);
         running.start();
    }
}
```

Figure 139 (Part 1 of 2). sign2 Class

```
public void stop(){
    if (running!=null){
       running.stop();
       running=null;
       }
}

public void run(){
 while(true)
       {
         d1=new Date();
         repaint();
         try{ Thread.sleep(10); }
         catch (InterruptedException e){ }
        }
}

public void paint(Graphics g) {
if (outImage != null){ // paint the image
       g.drawImage(outImage,0,0,null);
 }
}

public void update(Graphics g){
 Dimension dim=size();  // Take the Applet actual size
 //Verify the offscreen context
 if ((outGraphic == null)|| (dim.width != outDimension.width)
     || (dim.height != outDimension.height))
  {
   outDimension=dim;
   outImage=createImage(dim.width, dim.height);
   outGraphic=outImage.getGraphics();
  }  // And now prepare the outGraphics for the painting
 outGraphic.setColor(Color.white);
 outGraphic.fillRect(0,0,outDimension.width,outDimension.height);
 paintImage(outGraphic);
 g.drawImage(outImage,0,0,null); // put the out image in
}

public void paintImage (Graphics g)
{
 g.setFont(f);
 g.setColor(Color.blue);
 g.drawString(d1.toString(),5,20);
 g.drawString(text,x,40);
 if ((metrics.stringWidth(text)+x)==0)
  {
    x=size().width;
  }
 else
  {
    x=x-1;
  }
}
}
```

Figure 139 (Part 2 of 2). sign2 Class

In order to do the same with graphics in .gif files you can use the same code using the getImage() method.

In Figure 140 we found a small animator program (the SDK has an Animator class too) without sound. In Figure 143 on page 264 you are shown how to make Java work with sounds.

```
  import java.awt.*;

public class anim extends java.applet.Applet implements Runnable{
  Font f = new Font("TimesRoman",Font.BOLD,12);
  FontMetrics metrics = getFontMetrics(f);
  Thread running;
  String initial;
  String finish;
  String loop;
  String current;
  String baseName;
  int x;
  int i;
  Dimension outDimension;
  Graphics outGraphic;
  Image outImage;
  Image figures[];

public void init(){

   baseName=getParameter("base");
   initial=getParameter("initial");
   finish=getParameter("final");
   loop=getParameter("loop");
   current=new String(initial);
   if (loop==null) loop="yes";
   System.out.println("loop="+loop);
   if ((initial==null)||(finish==null)) destroy();

//retrieve the images before diplaying;
   x=difference(finish,initial);
   figures=new Image[x];
   for (i=0;i<x;i++){
     figures[i]=getImage(getCodeBase(),baseName+current+".gif");
     current=advance (current);
    }
//Beginning with the first figure;
   i=0;
}
```

Figure 140 (Part 1 of 3). Animation with Gif Files

```
public String advance(String string)
{
 int large=string.length();
 char last=string.charAt(large-1);
 StringBuffer buffer=new StringBuffer(string);
 last ++;
 buffer.setCharAt(large-1,last);
 string=buffer.toString();
 return string;
}

public int difference(String major, String minor){
 return (int)(major.charAt(major.length()-1)-minor.charAt(minor.length()-1));
}

public void start(){
   if (running==null) {
         running=new Thread(this);
         running.start();
    }
}

public void stop(){
    if (running!=null){
       running.stop();
       running=null;
       }
}

public void run(){
 boolean flag=true;
 while(flag)
       {
         repaint();
         if (i==x){
             if (loop.equals("no")){
                    flag=false;}
             else{
                        //Again the first picture
                    i=0;
                  }
            }
         try{ Thread.sleep(100); }
         catch (InterruptedException e){ }
        }
}

public void paint(Graphics g) {
if (outImage != null){ // paint the image
       g.drawImage(outImage,0,0,null);
 }
}
```

Figure 140 (Part 2 of 3). Animation with Gif Files

```
public void update(Graphics g){
 Dimension dim=size();  // Take the Applet actual size
 //Verify the offscreen context
 if ((outGraphic == null)|| (dim.width != outDimension.width)
     || (dim.height != outDimension.height))
  {
   outDimension=dim;
   outImage=createImage(dim.width, dim.height);
   outGraphic=outImage.getGraphics();
  }
// Prepare the outGraphics for the painting
 outGraphic.setColor(Color.white);
 outGraphic.fillRect(0,0,outDimension.width,outDimension.height);
 paintImage(outGraphic);
// put the out image in
 g.drawImage(outImage,0,0,null);
}

public void paintImage (Graphics g)
{
 g.setColor(Color.black);
 g.drawImage(figures[i],30,30,this);
 i++;
}

}
```

Figure 140 (Part 3 of 3). Animation with Gif Files

```
<HTML>
<Title>Sign Applet </Title>
<BODY bgcolor="#FFFFFF">
<APPLET Code="anim.class" WIDTH = 200 HEIGHT=160 ALIGN=RIGHT>
<PARAM Name=base Value="hello">
<PARAM Name=initial Value=1>
<PARAM Name=final Value="9">
The chip is not displayed because you are not using an Java anabled
Browser.
</Applet>
<Hr>
This is a good example for an animated applet.
</Body>
</HTML>
```

Figure 141. A Simple Animation Program

You can see in the init() method that there is a need to get all the parameters and begin to retrieve the images from the original place. The getCodeBase() method is returning the base directory of the URL where the applet is.

The advance method gives us a counter based on the last character of a String object. The Difference method tells us how many GIF files we are going to use.

The start() method and stop method are still the same. The run() method creates a loop on the images we have. Be careful with the string comparison. You can use the == operator to see if two strings are the same object, but if you have two different strings and you want to compare the content, you must

use the equals() method to do the comparison. If two objects are the same, it means they are sharing resources and memory space as well. If there are two strings that are not sharing resources but the content is equal for both of them, you are not able to know it with the == operator; the same happens with your objects. If you need to create comparisons between two objects you have to program their methods.

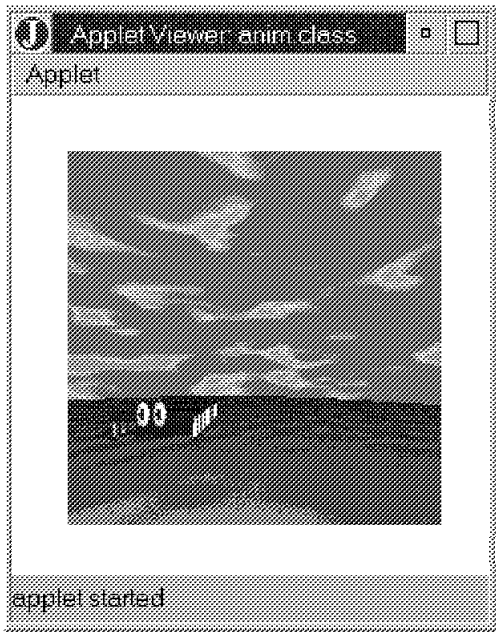


Figure 142. The anim Class with the holax.gif Files

The formats that Java accepts with the getimage() method are JPEG (.JPG) and CompuServe (.GIF).

The use of getImage() and getAudioClip is very similar.

The syntax for both of them is:

```
String URLstring;
  .
  .
getImage(URLCodeBase,File);
getAudioClip(URLCodeBase,File);
```

In both of them the parameters are strings indicating the directory (URLCodeBase) where the image or the sound is and the name of the file inside the directory. The getImage() method will return an Image object and the getAudioClip() returns an AudioClip object. The AudioClip objects has the following methods:

```
loop()
play()
stop()
```

As long as you can guess, the loop begins to play the audio and creates a loop playing it when the AudioClip gets the final. The play() begins to play from the beginning. You can see the example for the use of the Audioclip object in the animation on Figure 143 on page 264.

```
  import java.awt.*;

public class anim extends java.applet.Applet implements Runnable{
  AudioClip audio;
  Thread running;
  String initial;
  String finish;
  String loop;
  String current;
  String baseName;
  String audioName;
  int x;
  int i;
  Dimension outDimension;
  Graphics outGraphic;
  Image outImage;
  Image figures[];

public void init(){
   audioName=getParameter("audio");
   baseName=getParameter("base");
   initial=getParameter("initial");
   finish=getParameter("final");
   loop=getParameter("loop");
   current=new String(initial);
   if (audioName!=null)
     {
      audio=getAudioClip(getCodeBase(),baseName+audioName+".au");
     }
   if (loop==null) loop="yes";
   System.out.println("loop="+loop);
   if ((initial==null)||(finish==null)) destroy();

//retrieve the images before diplaying;
   x=difference(finish,initial);
   figures=new Image[x];
   for (i=0;i<x;i++){
     figures[i]=getImage(getCodeBase(),baseName+current+".gif");
     current=advance (current);
    }

//Beginning with the first figure;
   i=0;
}
```

Figure 143 (Part 1 of 3). Animation with Sound

```
public String advance(String string)
{
 int large=string.length();
 char last=string.charAt(large-1);
 StringBuffer buffer=new StringBuffer(string);
 last ++;
 buffer.setCharAt(large-1,last);
 string=buffer.toString();
 return string;
}

public int difference(String major, String minor){
 return (int)(major.charAt(major.length()-1)-minor.charAt(minor.length()-1));
}

public void start(){
   if (running==null) {
         running=new Thread(this);
         if (audio!=null) audio.loop();
         running.start();
    }
}

public void stop(){
    if (running!=null){
       running.stop();
       if (audio!=null) audio.stop();
       running=null;
       }
}

public void run(){
 boolean flag=true;
 while(flag)
       {
         repaint();
         if (i==x){
             if (loop.equals("no")){
                    flag=false;}
             else{
                        //Again the first picture
                   i=0;
                  }
            }
         try{ Thread.sleep(100); }
         catch (InterruptedException e){ }
        }
}

public void paint(Graphics g) {
if (outImage != null){ // paint the image
       g.drawImage(outImage,0,0,null);
 }
}
```

Figure 143 (Part 2 of 3). Animation with Sound

```
public void update(Graphics g){
 Dimension dim=size();  // Take the Applet actual size
 //Verify the offscreen context
 if ((outGraphic == null)|| (dim.width != outDimension.width)
     || (dim.height != outDimension.height))
  {
   outDimension=dim;
   outImage=createImage(dim.width, dim.height);
   outGraphic=outImage.getGraphics();
  }
// Prepare the outGraphics for the painting
 outGraphic.setColor(Color.white);
 outGraphic.fillRect(0,0,outDimension.width,outDimension.height);
 paintImage(outGraphic);
// put the out image in
 g.drawImage(outImage,0,0,null);
}

public void paintImage (Graphics g)
{
 g.setColor(Color.black);
 g.drawImage(figures[i],30,30,this);
 i++;
}

}
```

Figure 143 (Part 3 of 3). Animation with Sound

Some methods that can help us to find the right URL to call are:

getCodeBase(); //retrieves the directory in which the applet is

getDocumentBase() // Gets the document URL

5.15 Events Handling

When you are programming an interactive applet you may want to make some responses to the user's actions, such as clicking on some areas and creating buttons or text boxes like you have done in some applications.

All of those objects are part of the awt package. A package is a group of classes that are in the same category or were compiled from the same file. A package can be useful to create more than one class in a single file or to create long programs with multiple classes on it. The awt package contains all of the Windows classes that are useful and commonly used. When someone clicks on or writes something in a text box, an event is driven. For the programmer, the most important thing to know is who the receptor of the event is. To do this you must override the handleEvent(Event e) method; this is a public boolean method.

The Event class has one property named target that contains the object that was the target of the current event.

The most useful property on the event handling method is the ID property on the Event object. With this property you can know the type of event to handle.

5.15.1 The Mouse Event Handler

When you do not override the handleEvent(Event e) some events are handled by the default method making a call to other methods. Some of them are responses to the mouse action, but the methods you have for these actions are empty. All this means that when your mouse moves, or you click in some place of the applet, you call an empty method; you can say you are doing nothing.

To create an action when an event is called, you override these methods, the mouse and keyboard methods are part of the default methods that already exist. The functions you override for the mouse are:

```
public boolean mouseUp(Event e, int x, int y)
public boolean mouseDown(Event e, int x, int y)
public boolean mouseDrag(Event e, int x, int y)
public boolean mouseMove(Event e, int x, int y)
public boolean mouseEnter(Event e, int x, int y)
public boolean mouseExit(Event e, int x, int y)
```

The first two methods can be used for the actual mouse click. The actions required here can be coded in there; the method should return true, only to create an acknowledgement. The mouse event can or cannot occur in the same place (you can have a Drag event going on), and the actions taken in the up or down can be completely different. One good exercise can be to take the anim program and make it stop or restart with a click on the applet.

The mouseDrag(Event e, int x, int y) occurs in all the points where the dragging is done. A good example of the dragging method can be a hand free drawing program. We are going to need some special objects from the Graphics object, the Point object, the Event object and the Color object. The self explanatory program shown in Figure 144 does the free hand drawing.

```
import java.awt.Graphics;
import java.awt.Color;
import java.awt.Event;
import java.awt.Point;
public class freeHand extends java.applet.Applet{
 Point mypoint;
 public void init(){
    setBackground(Color.white);
    }/*end of init*/
 public boolean mouseDrag(Event e, int x, int y){
   mypoint=new Point(x,y);
   repaint();
   return true;
   }/*end of mouseDrag*/
 public void update(Graphics g){
   g.setColor(Color.black);
   g.fillOval(mypoint.x,mypoint.y,3,3);
  }
}/*end of class*/
```

Figure 144. freeHand Object. The freeHand object is a good explanation for the use of the mouse methods.

5.16 AWT (Abstract Window Toolkit)

Here we are make a brief description of some components of the AWT, but you can find a complete reference in the Java Tutorial Home Page:

http://java.sun.com/tutorial

You realize how dirty the freeHand applet can get. One of the simplest solutions is to create a button to clean up all of the mess. A simple button can solve the problem. Clean the window when the button is clicked, drawing a filled rectangle with the same size as the applet size.

A button is part of the AWT package and is one of the main components. The components that Java has in the Abstract Window Toolkit are:

Button	Canvas	CheckBox	CheckBoxGroup	Choice
Dialog	FileDialog	Frame	Label	List
Menu	MenuBar	MenuItem	Panel	Scrollbar
TextArea				

All of these objects generate a different class of messages, and these messages are captured like events in the public boolean handleEvent(Event e) method.

To add any of these components you use the add() method in the following way:

```
Button buttonOk;
public void init(){
  buttonOk = new button("Ok");
  add(button(Ok);
}
```

Other useful objects can be the labels; you can add them the same way you did before. The constructors for the label are:

```
Label()
Label(String)
Label(String, alignment)
```

The alignment is an integer, but you can use the following variables to make the alignment easier: Label.RIGHT, Label.LEFT and Label.CENTER.

Remember there are no constant variables in Java; the approach to a constant can be the final word.

The buttons are placed in the most convenient place defined for the applet, but you can use layout managers to put the button in the most convenient place for you; we'll get back to this later.

Other important components are the check boxes, radio buttons, choice menus and text fields. The check boxes are interfaces with two possible values for each one, true or false, and they can be used in exclusive or nonexclusive ways. When you use the exclusive way you can only select one of them; when you are using the nonexclusive mode you can select more than one check box at a time.

You cannot group the check boxes, but you can group the radio buttons, that work alike but only one radio button in the group can be selected. To create a radio button group, you must create a new CheckboxGroup object and add new check boxes to them. When the check boxes are grouped, they are converted to radio buttons.

Another type of control is the choice menu. This object creates a pull-down menu, where you can choose an item. To add an item to the choice menu you must create a Choice object and use the addItem(String) method to add a new item to the list. In Figure 145 on page 272 you can see a program that creates all the objects.

Other kinds of components useful for an interface can be the text fields and text areas. Both of them can get input from the user, but in the text area you can get a multiline response and put scrollbars in it.

The scrollbars and sliders are useful to help you manipulate values when the user wants to. You must use the Scrollbar class to represent both of them.

The Canvas class allows you to put a graphics container where you can draw or put an image on it, but you cannot add other components such as buttons or text fields.

Table 24 (Page 1 of 2). Constructors for the Given AWT Components

Object	Constructors	Explanation
Label	Label()	Creates a new label without text and left justified.
	Label(String label)	Creates a new label containing the given string, left justified.
	Label(string label, int alignment)	Creates a label containing the given string and with the given alignment; you can use Label.LEFT, Label.CENTER or Label.RIGHT to indicate the alignment.
Button	Button()	Creates a new button, no text on it.
	Button(String label)	Creates a new button labeled with the given text.
Checkbox	Checkbox()	Creates a check box that is not a part of any CheckboxGroup and has no label.
	Checkbox(String label)	Creates a new check box with a label containing the given string.
	Checkbox(String label, CheckboxGroup group, boolean state)	Creates a new check box with a label, this one is part of the given CheckboxGroup and his initial state is also given. You can use null to indicate that the check box is not part of any group.
CheckboxGroup	CheckboxGroup()	Creates a CheckboxGroup, a special check box container.
Choice	Choice()	Creates a Choice menu.
TextField	TextField()	Creates a new TextField, empty.
	TextField(int cols)	Creates a new TextField that is cols characters width.
	TextField(String contains)	Creates a new TextField with a default text.
	TextField(String Text, int cols)	Creates a TextField with a default text and cols character width.
TextArea	TextArea()	Creates a new TextArea object.
	TextArea(int rows, int cols)	Creates a TextArea object rows lines height and cols characters width.
	TextArea(String text)	Creates a TextArea containing the specified text.
	TextArea(String text, int rows, int cols)	It is a combination of the other constructors.

Table 24 (Page 2 of 2). Constructors for the Given AWT Components		
Object	**Constructors**	**Explanation**
List	List()	Creates a new scroll list object.
	List(int rows, boolean multiple)	Creates a new scroll list with a number of visible rows. The state is used to indicate if the list can have multiple selections on it.
Scrollbar	Scrollbar()	Creates a new vertical scrollbar.
	Scrollbar(int orientation)	Creates a new scrollbar, the orientation can be Scrollbar.VERTICAL or Scrollbar.HORIZONTAL
	Scrollbar(int orientation, int value, int visible, int minimum, int maximum)	Creates a new scrollbar, the orientation is used as the above constructor, you must indicate the minimum and maximum values of the scrollball. The visible parameter is used to say the size represented by the bubble in the scrollbar.
Canvas	Canvas()	Create a new canvas.

Constructors and methods are listed in these two tables. Some methods are not listed but are useful. For example, the method addItem in the choice menu objects or in the scroll list can be used to add items on them. For a complete reference of the API, go to the following URL:

`http://java.sun.com/products/JDK/CurrentRelease/api`

Table 25 (Page 1 of 2). Some Methods of the AWT Components		
Object	**Method**	**Action**
Label	getText()	Returns a string containing this label's text.
	setText(String)	Changes the text of this label.
	getAlignment()	Returns an integer representing the alignment of this label: `Ø is Label.LEFT, 1 is Label.CENTER, 2 is Label.RIGHT`
	setAlignment (int)	Changes the alignment of this label to the given integer, use the class variables above.
Checkbox	getLabel()	Returns the string that is contained on the label.
	setLabel(String)	Set a new text on the check box label.
	getState()	True or false. If it is selected the method returns true.
	setState(boolean)	Changes the state of a check box by the specified value on parameter.
Choice	getItem(int)	Returns the string of the item in the specified position.
	countItems()	Returns the number of items on the choice menu.
	getSelectedIndex()	Returns the current selected item position.
	getSelectedItem()	Returns the current item string.
	select(int)	Selects the item in the given position.
	select(String)	Selects the item with the given string.

Table 25 (Page 2 of 2). Some Methods of the AWT Components

Object	Method	Action
Text Field	getText()	Returns the text field contains.
	setText(String)	Changes the text.
	getColumns()	Returns the width of the TextField object.
	select(int, int)	Selects the text between the two given positions.
	selectAll()	Selects all the text.
	isEditable(boolean)	Returns the state of the TextField, true if it is enabled.
	setEditable(boolean)	Enables the TextField object, allowing the user to edit it.
	getEchoChar()	Returns the mask input character.
	echoCharIsSet()	Returns true if a masking character was set.
Text Area (most of TextField applies)	getColumns()	Returns the width of the text area in character columns.
	getRows()	Returns the number of rows that the text area has.
	insertText(String, int)	Insert a string at given position. Remember the first position on every string is 0.
	replaceText(string text, int beginning, int end)	Replaces the text between beginning and end with the specified text.
	setLineIncrement(int inc)	Sets the number of rows of movement when the inside part of the scrollbar is clicked, the default value is 10.
	getLineIncrement()	Returns the value of movement when the inside part of the scrollbar is clicked.
	setPageIncrement(int inc)	The same as above, but it sets the vertical movement of the text area.
	getPageIncrement()	The same as getLineIncrement, but with the vertical movement instead.
List	getItem(int)	Returns the string in the given position.
	countItems()	Returns the number of items on the object.
	getSelectedIndex()	Returns the item selection. Valid for single selections only.
	getSelectedIndexes()	Returns an array of the selected positions.
	getSelectedItem()	Returns the selected item as string. Single selections only.
	getSelectedItems()	Returns an array of strings with the selected items.
	select(int)	Selects the given position.
	select(String)	Selects the item with that string.
Scrollbar	getMaximum()	Returns the maximum value of the scrollbar.
	getMinimum()	Returns the minimum value of the scrollbar.
	getOrientation()	Returns the orientation of the scrollbar.
	getValue()	Returns the current value of the scrollbar.
	setValue(int)	Sets a new value for the scrollbar.

The next program shows how to make use of the AWT of Java and capture the button event. Knowing the string of the button, you can control the correct action for an applet or window on your class.

```
import java.awt.*;

public class awtexample extends java.applet.Applet{
CheckboxGroup firstGroup;
Choice cho;
List scrolllist;
public void init(){
        cho=new Choice();
        setBackground(Color.white);
        firstGroup = new CheckboxGroup();
        add (new Checkbox("Hello", firstGroup, false));
        add (new Checkbox("Bye Bye", firstGroup, true));
        add (new Button("OK"));
        cho.addItem("Monkey");
        cho.addItem("Frog");
        cho.addItem("Bull");
        add(cho);
        scrolllist=new List(3,true);
        scrolllist.addItem("Carlos de Luna");
        scrolllist.addItem("Patrick Schmitt");
        scrolllist.addItem("Nilson Batista");
        scrolllist.addItem("Marcio Venzi");
        scrolllist.addItem("Roberto Oku");
        add(scrolllist);
}
public boolean action (Event e, Object arg){
if (e.target instanceof Button)
    System.out.println("A button was pressed");
return true;
}
}
```

Figure 145. AWT Components Example

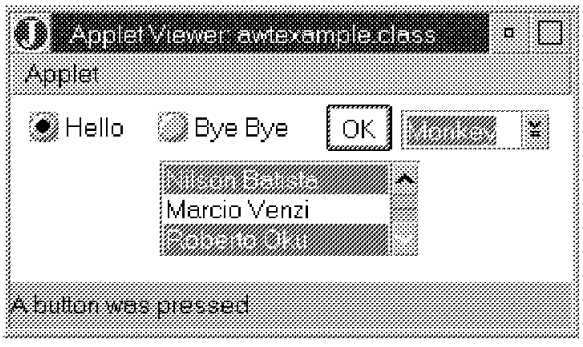


Figure 146. AWT Components on OS/2

5.16.1 Layouts and Panels

In order to control the place where the buttons and all other components are going to be placed, you can use the panels and layouts. Layouts are a kind of presentation that can be used to create better interfaces with the Java AWT. The layouts are special containers that can calculate the right coordinates where you must put the buttons to get the right presentation.

The most common layouts are:

1. The FlowLayout is used to arrange buttons in the panel; this is the default layout manager for the applet class.
2. The GridLayout is used to have a rectangular grid. The container is split into equal-sized rectangles and every component is placed on one rectangle.
3. The GridBagLayout is used like the most flexible layout of all. It aligns the component vertically and horizontally without requiring that the components be the same size.
4. The BorderLayout allows you to put 5 components using the North, South, East, West and Center strings.
5. CardLayout allows you to contain several cards in the same container, but only one is visible.
6. The Insets allow you to leave an inset on each side of the screen.

To use a layout you must use the setLayout() method object like a parameter.

You can find examples and an explanation on the use of the layouts by going to the following URL:

`http://java.sun.com/tutorial`

Or, look at the API specifications to see how to use them and a complete reference for the functions.

5.17 URL Management

The URL management is part of the java.net package. This packages allows the programmer to use connections using streams and UNIX-like sockets. It also allows you to use the HTTP protocol to retrieve files. Using these functions you can create stand-alone applications or you may use it to create better pages.

At the beginning of this chapter we explained the applet restrictions, but you can still do good things with them.

In this part of the chapter we are covering only how to create links from a Java applet; you can look in the tutorial or the API home pages for a complete reference.

To create a new connection you must use a URL object. This object represents the Uniform Resource Locator. Its components are the protocol (http, ftp, gopher, news, etc.), the node (www.ibm.com, www.mexico.ibm.com, java.sun.com, etc.), a connection port (if you use http the well known port is 80, some sites can use other ports) and the file. Sometimes the URL also contains an anchor (http//.../filename#anchor).

The URL object allows you to make connections only to retrieve and show (make links) to other pages. Other types of connections must be with the same server that owns the applet. To create a connection, you must use streams and control them (refer to the Java Tutorial, the Java API and other publications mentioned in the bibliography).

The next step after you have the URL ready is to show the page. To do this you use the getAppletContext().showDocument(URL Document) method. In the example shown in Figure 147 you can see the use of layouts and buttons to connect different sites.

```
import java.awt.*;
import java.net.URL;
import java.net.MalformedURLException;
public class conURL extends java.applet.Applet
{
 References URLlist[]= new References[9];
 public void init()
   {
    URLlist[0]=new References("Charly's","http://www.cem.itesm.mx/cluna/mio.html");
    URLlist[1]=new References("ITSO","http://www.redbooks.ibm.com");
    URLlist[2]=new References("IBM JAVA","http://www.hursley.ibm.com/javainfo");
    URLlist[3]=new References("SUN Java","http://java.sun.com");
    URLlist[4]=new References("Java Tutorial","http://java.sun.com/tutorial");
    URLlist[5]=new References("Java API","http://java.sun.com/products/JDK/CurrentRelease/api");
    URLlist[6]=new References("IBM","http://www.ibm.com");
    URLlist[7]=new References("IBM Mexico","http://www.mexico.ibm.com");
    URLlist[8]=new References("IBM Brazil","http://www.ibm.com.br");
    setLayout(new GridLayout(3,3,5,5));
    for (int i=0; i<9;i++)
    {
      add(new Button(URLlist[i].Name));
    }
   }/*endinit*/
```

Figure 147 (Part 1 of 2). URL Example

```
 public boolean action (Event e, Object where)
 {
  if (e.target instanceof Button)   /*A button was clicked*/
  {   moveto((String)where);    /*method to call the URL*/
      return true;}
  else
      return false;
 }/*end action*/
 public void moveto(String where)
 {
  boolean flag=false;
  URL auxURL=null;
  for (int i=0; (i<9)&&(!flag);i++)
   {
    if(where.equals(URLlist[i].Name)){
     auxURL=URLlist[i].theURL;
     flag=true;
    }
   }/*endfor*/
   if (auxURL !=null){
    getAppletContext().showDocument(auxURL);
   }
  else
  {
   System.out.println("The choosen reference is not a URL");
  }
 }/*endmoveto*/
}/*end class*/
class References{
    String Name;
    String Reference;
    URL theURL;
References (String NAME, String REFERENCE){
  this.Name=new String(NAME);
  this.Reference=new String(REFERENCE);
  try{ this.theURL=new URL(REFERENCE);
  }catch(MalformedURLException e){
   this.theURL= null;
   System.out.println("This is not a URL reference");
  }
 }
}
```

Figure 147 (Part 2 of 2). URL Example

This program will give you a good idea of how to do connections and links from your own page. Adding additional windows to your applets could be a good idea. The applets windows will always have the sign Warning: this is an applet window. To do this you must use the Frame class. For more information you can look for the following biblography:

Java Tutorial, by Marie Campione and Cathy Walrath, part of The Java Series, publishedby Addison Wesley.

http://java.sun.com/tutorial

Teach Yourself Java in 21 Days by Laura Lemay and Charles L. Perkins Published by Sams.net Publishing, Indianapolis, IN, USA.

Java in a Nutshell, by David Flannagan, published by O'Reilly.

Hooked on Java, by Gary Cornell and Cary Horstman Published by Addison Wesley.

5.18 Brief Guide to Advanced Topics

Java programming allows you to have a multiserver machine, creating classes that can communicate with this server using the applets. Powerful classes can be developed to enhace your existing server applications, making a common interface using the Web browsers like viewers.

Some topics you must review to get the most of Java are:

1. Packages: They allow you to create huge programs and hide classes or have more than one public class on the same package.
2. The java.net package: This is the tool you use to create networking applications and create secure applications on the net.
3. Review your C and C + + knowledge to create native interfaces with Java, but remember, if you do this your applications will not be platform independent.

These areas will help you to improve your applications and get the most out of the Internet, enjoy your programming and create a good job. Remember to see the bibliography for information about Java and the themes listed above. Some useful hints to create applets are:

1. Always override the public String getAppletInfo() returning your copyright information.
2. Take out all the System.out.println lines; if you want to put something for the knowledge of the user you must use the showStatus() method.
3. Always implement the stop() and run() methods if you are generating graphics or multithreading applets.
4. It is a very good idea to give a Form to the sound tracks to stop (a button, a simple click on the applet, anything). Be kind to your visitors (if you want to have them back).
5. Remember, the more flexible your program the more helpful it is.
6. Enjoy your programming. If you do, the users will note it is a good job.

More information on how to do better and larger programs is in the bibliography above. If you want to see something special on the following versions contact us at IBM by filling out the form at the end of this book or e-mail: cdeluna@vnet.ibm.com.

5.19 When to Consider CGI and When to Consider JAVA

If you already read this chapter and the one that talks about CGI, you can answer the question yourself. CGIs are a good tool to make an interface when you need to store data in your server, create some special processing that represents a large program or just to get information about the client (visitor). Java is a good tool to create interactive pages, let the client make their own

calculations, and have small programs that allow you to communicate with the server for a better interaction. Java is a language to do stand-alone applications as well; it has all the advantages of the object-oriented languages and it is binary portable across all the platforms; creating stand-alone client/server applications with Java can be a very good idea.

Other tools such as JavaScript can be useful to create interactive pages when you have forms or you want to make applications, but only a few browsers support it. It is not as powerful as Java and the philosophy of it is only a help to the HTML language. If you have an application that can be inside of a form, you can think about JavaScript. If you need something more than an interactive form or you want to do it compatible with most of the browsers, you should use Java.

Chapter 6. Multimedia Concepts and Terms

This chapter gives you an overview of the multimedia concepts and terms used in the Internet environment. The following are common image formats on the Internet.

6.1.1 JPEG Image Format

JPEG (pronounced jay-peg) is a standardized image compression mechanism. JPEG stands for Joint Photographic Experts Group, the original name of the committee that wrote the standard. All graphical browsers support the JPEG format. JPEG is designed for compressing either full-color or gray-scale images of natural, real-world scenes. It works well on photographs, naturalistic artwork, and similar material, but not so well on lettering, simple cartoons, or line drawings.

JPEG handles only still images, but there is a related standard called MPEG for motion pictures. JPEG is *lossy*, meaning that the decompressed image isn′t quite the same as the one with which you started. There are lossless image compression algorithms, but JPEG achieves much greater compression than is possible with lossless methods.

JPEG is designed to exploit known limitations of the human eye, notably the fact that small color changes are perceived less accurately than small changes in brightness. Thus, JPEG is intended for compressing images that will be looked at by humans. If you plan to machine-analyze your images, the small errors introduced by JPEG may be a problem for you, even if they are invisible to the eye.

A useful property of JPEG is that the degree of lossiness (loss resolution) can be varied by adjusting compression parameters. This means that the image maker can trade off file size against output image quality. You can make extremely small files if you don′t mind poor quality; this is useful for applications such as indexing image archives. Conversely, if you aren′t happy with the output quality at the default compression setting, you can jack up the quality until you are satisfied and accept lesser compression.

Another important aspect of JPEG is that decoders can trade off decoding speed against image quality by using fast but inaccurate approximations to the required calculations. Some viewers obtain remarkable speedups in this way. There are two good reasons to use JPEG against other formats: to make your image files smaller, and to store 24-bit-per-pixel color data instead of 8-bit-per-pixel data.

Making image files smaller is a win for transmitting files across networks and for archiving libraries of images. Being able to compress a 2-MB full-color file down to, for example, 100 KB makes a big difference in disk space and transmission time. JPEG can easily provide 20:1 compression of full-color data. If you are comparing GIF and JPEG, the size ratio is usually more like 4:1.

If your viewing software doesn′t support JPEG directly, you′ll have to convert JPEG to some other format to view the image. Even with a JPEG-capable viewer, it takes longer to decode and view a JPEG image than to view an image of a simpler format such as GIF. Thus, using JPEG is essentially a time/space tradeoff: you give up some time in order to store or transmit an image more

© Copyright IBM Corp. 1996

cheaply. But it's worth noting that when network or telephone transmission is involved, the time savings from transferring a shorter file can be greater than the time needed to decompress the file.

The second fundamental advantage of JPEG is that it stores full color information: 24 bits/pixel (16 million colors). GIF, the other image format widely used on the net, can only store 8 bits/pixel (256 or fewer colors). GIF is reasonably well matched to inexpensive computer displays. Most run-of-the-mill PCs can display no more than 256 distinct colors at once. But full-color hardware is getting cheaper all the time, and JPEG images look much better than GIFs on such hardware. Within a couple of years, GIF will probably seem as obsolete as black-an d-white MacPaint format does today. Furthermore, JPEG is far more useful than GIF for exchanging images among people with widely varying display hardware, because it avoids prejudging how many colors to use. Hence, JPEG is considerably more appropriate than GIF for use as a Usenet and World Wide Web standard format.

Many people are scared off by the term *lossy compression*. But when it comes to representing real-world scenes, no digital image format can retain all the information that impinges on your eyeball. By comparison with the real-world scene, JPEG loses far less information than GIF. The real disadvantage of lossy compression is that if you repeatedly compress and decompress an image, you lose a little quality each time.

JPEG does not support transparency and is not likely to do so any time soon. It turns out that adding transparency to JPEG would not be a simple task. The traditional approach to transparency, as found in GIF and some other file formats, is to choose one otherwise-unused color value to denote a transparent pixel. That can't work in JPEG because JPEG is lossy: a pixel won't necessarily come out the exact same color that it started. Normally, a small error in a pixel value is OK because it affects the image only slightly. But if it changes the pixel from transparent to normal or vice versa, the error would be highly visible and annoying, especially if the actual background were quite different from the transparent color.

A more reasonable approach is to store an alpha channel (transparency percentage) as a separate color component in a JPEG image. That could work since a small error in alpha makes only a small difference in the result. The problem is that a typical alpha channel is exactly the sort of image that JPEG does very badly on: lots of large flat areas and sudden jumps. You'd have to use a very high quality setting for the alpha channel. It could be done, but the penalty in file size is large. A transparent JPEG done this way could easily be double the size of a non-transparent JPEG. That's too high a price to pay for most uses of transparency.

The only real solution is to combine lossy JPEG storage of the image with lossless storage of a transparency mask using some other algorithm. Developing, standardizing, and popularizing a file format capable of doing that is not a small task. and transparency doesn't seem worth that much effort.

6.1.2 GIF Image Format

The GIF image format uses a built-in LZW compression algorithm. This compression algorithm is patented technology and currently owned by Unisys Corporation. As of 1995, Unisys decided that commercial vendors, whose products use the GIF LZW compression, must license its use from Unisys. End users, online services, and non-profit organizations do not pay this royalty. Since it's inception, GIF has been a royalty-free format. Only as of 1995 did Unisys decide to collect royalties. To avoid this royalty, vendors have developed an alternative to GIF that supports transparency and interlacing called PNG (ping), the Portable Network Graphic. To our knowledge PNG, however, does not support a multiple image data stream.

The GIF87a allowed for the following features:

- LZW compressed images
- Multiple images encoded within a single file
- Positioning of the images on a logical screen area
- Interlacing

This means that nine years ago it was possible to do simple animation with GIFs by encoding multiple images, what we will refer to as frames, in a single file. GIF89a is an extension of the 87a spec. GIF89a added:

- How many 100ths of a second to wait before displaying the next frame
- Wait for user input
- Specify transparent color
- Include unprintable comments
- Display lines of text
- Indicate how the frame should be removed after it has been displayed
- Application-specific extensions encoded inside the file

Netscape Navigator is the only browser than comes close to full GIF89a compliance. The lines of text and user input are not currently supported in Navigator 2.0, and the image removal doesn't support removal by the previous image. Most browsers support single image GIF87a and will only recognize the transparency flag of GIF89a.

GIF89a is still a 256-color (maximum) format. GIF allows for any number of colors between 2 and 256. The fewer the colors the less data and the smaller the graphic files. If your GIF only uses 4 colors, you can reduce the palette to only 2 bits (4 color) and decrease the file size by upwards of 75%.

The following software lets you set bits-per-pixel for GIFs:

- Adobe Photoshop
- Fractal Painter
- Painter 2.0
- PhotoStudio
- PhotoGIF
- PaintShop Pro

- PaintIt
- WebImage

GIFs are composed of Blocks and Extensions. Blocks can be classified into three groups:

- Control
- Graphic-Rendering
- Special Purpose

Control blocks, such as the Header, the Logical Screen Descriptor, the Graphic Control Extension and the Trailer, control how the graphic data is handled. Graphic-rendering blocks such as the Image Descriptor and the Plain Text Extension contain data used to render a graphic . Special purpose blocks such as the Comment Extension and the Application Extension are not used by GIF decoders at all. The Logical Screen Descriptor and the Global Color Table affect all the images in a single file. Each Control block will only affect a single Image block that immediately follows it. A GIF file contains a global palette of common colors for all the images in its file to work from. This palette can have 2, 4, 8, 16, 32, 64, 128, or 256 defined colors. Palettes are very important. Every color displayed in your GIF must come from a palette. The fewer colors used, the easier it will be for systems to display your images. The global palette is applied to all images in a GIF file. If an individual image differs greatly from that global palette, it may have a local palette that affects its color *only*. However, no image can every reference more than one palette, so 256 colors per image is the max. Having a bunch of local palettes with wildly varied colors can sometimes cause color shifts in your display.

The following are the benefits of using GIF images:

- All the benefits of GIF: transparency, compression, interlacing, 2, 4, 8, 16, 32, 64, 128 and 256 color palettes for optimum size and compression.
- Supported by the basic Netscape product and no plug-ins or additional software. Tested on Win 3.1x, Win95, MAC, UNIX, Sun, Linux, and Irix.
- Web designer does not need access to Internet provider′s web server, server-side includes (SSI), or CGI/PERL scripting. If you have a program that can make multi-image 89a GIFs, you can make this animation.
- The animation is repeatable and reusable. You can place the same image on a page multiple times. It performs a single download for all and loops all from the cache.
- The animation only loads once, so your modem doesn′t keep downloading constantly. It is faster than server-reliant methods.
- The animations are surprisingly compact.
- Anyone can use them on their page. Anyone with a web page can include this animation. In fact, if you save any of the animated GIFs to your hard drive, you will have the entire animation to put in your own pages. Please contact the creator for usage.
- Works like any other GIF; include on your page in an IMG or FIG tag, even anchor it; it works invisibly.

The following are the limitations of using GIF:

- All the limitations of GIFs: max of 256 colors, photographs are better compressed by JPEG.
- Only plays in Netscape 2.0 or higher, but does work with many platforms (Windows, MAC, UNIX, etc.).
- Will play once or continuously. Refresh will not play the image again, but reload or resizing the windows will. If the viewer returns back to the page from elsewhere, the image will play, even if cached. Later revisions of Navigator may support finite iterations of the animations.
- It cannot be used as a background GIF. Only the first frame will display.

Compuserve released the technical specification for GIF89a in July of 1989. The technical specification is an exact breakdown of the byte-for-byte structure and rules for interpreting and building this format.

6.2 Audio File Formats

Historically, almost every type of machine used its own file format for audio data, but some file formats are more generally applicable. In general, it is possible to define conversions between almost any pair of file formats. However, sometimes you lose information.

File formats are a separate issue from device characteristics. There are two types of file formats: self-describing formats, where the device parameters and encoding are made explicit in some form of header, and *raw* formats, where the device parameters and encoding are fixed.

Self-describing file formats generally define a family of data encodings, where a header field indicates the particular encoding variant used. Headerless formats define a single encoding and usually allow no variation in device parameters (except sometimes sampling rate, which can be a pain to figure out other than by listening to the sample). The header of self-describing formats contains the parameters of the sampling device and sometimes other information (for example, a human-readable description of the sound, or a copyright notice).

Most headers begin with a simple *magic word*. Some formats do not simply define a header format, but may contain chunks of data intermingled with chunks of encoding information. The data encoding defines how the actual samples are stored in the file (for example, signed or unsigned, as bytes or short integers, in little-endian or big-endian byte order, etc.). Strictly spoken, channel interleaving is also part of the encoding, although so far I have seen little variation in this area. Some file formats apply some kind of compression to the data (for example, Huffman encoding or simple silence deletion).

Here's an overview of popular file formats.

Extension, name	Origin	Variable parameters
au or snd	NeXT, Sun	rate, #channels, encoding, info string
aif(f), AIFF	Apple, SGI	rate, #channels, sample width, lots of info
aif(f), AIFC	Apple, SGI	same (extension of AIFF with compression)
iff, IFF/8SX	Amiga	rate, #channels, instrument info (8 bits)
voc	Soundblaster	rate (8 bits/1 ch; can use silence deletion)
wav, WAVE	Microsoft	rate, #channels, sample width, lots of info
sf	IRCAM	rate, #channels, encoding, info

Extension, name	Origin	Variable parameters
none, HCOM	Mac	rate (8 bits/1 ch; uses Huffman compression)
mod or nst	Amiga	(see below)

Note that the filename extension .snd is ambiguous; it can be either the self-describing NeXT format or the headerless Mac/PC format, or even a headerless Amiga format.

IFF/8SVX allows for amplitude contours for sounds (attack/decay/etc). Compression is optional (and extensible) and volume (author, notes and copyright properties, etc.) is variable.

AIFF, AIFC and WAVE are similar in spirit but allow more freedom in encoding style (other than 8 bit/sample), amongst others.

There are other sound formats in use on Amiga by digitizers and music programs, such as IFF/SMUS.

DEC systems use a variant of the NeXT format that uses little-endian encoding and has a different number.

Standard file formats used in the CD-I world are IFF, but on the disc they are in real-time files.

An interesting *interchange format* for audio data is described in the proposed Internet Standard MIME, which describes a family of transport encodings and structuring devices for electronic mail. This is an extensible format, and initially standardizes a type of audio data dubbed audio/basic, which is 8-bit U-LAW data sampled at 8000 samples/sec.

Finally, a somewhat different but popular format are MOD files, usually with the extension .mod or .nst (they can also have a prefix of mod.). This originated at the Amiga but players now exist for many platforms. MOD files are music files containing two parts:

1. A bank of digitized samples
2. A sequencing information describing how and when to play the samples

6.3 Musical Instruments Digital Interface (MIDI)

This international standard for digital music was established in 1982. It specifies the cabling and hardware required for connecting electronic musical instruments and computer systems. MIDI also specifies a communication protocol for passing data from one MIDI device to another. Any musical instrument can become a MIDI device by having the correct hardware interfaces and MIDI messages processing capabilities. Devices communicate with each other by sending messages that are digital representations of a musical score. MIDI data may include items such as sequences of notes, timings, instrument designations and volume settings. The standard multimedia platform can play MIDI files through either internal or external synthesizers. External MIDI devices are connected to the computer via the sound card′s MIDI port. MIDI expands the audio options available when developing multimedia. Use of MIDI is attractive because MIDI files require minimal storage space compared to digitized audio files, such as .WAV files.

MIDI ports are used to send and receive MIDI data. There can be many MIDI ports installed in a system. Each MIDI port contains a MIDI IN, MIDI OUT, and MIDI THRU connection. MIDI IN receives messages sent from other MIDI devices. MIDI OUT transmits messages that are originating from the local device to other MIDI systems. MIDI THRU forwards messages that were received by the MIDI IN to other devices. Each port can handle 16 MIDI channels. A synthesizer is the device which produces sound. Generally it has a built-in keyboard. There are several different methods used in synthesizer technology to produce music instrument sounds. By altering standard wave forms, such as the sign wave, a variety of sounds can be produced. Another method of producing sound is by playing back stored samples of real instruments. The newest synthesizer technology employs powerful computer technology to emulate musical instruments via mathematical algorithms that represent certain aspects of an instrument (for example, a bowed string, pipe blown). This technology gives musicians the ability to play a realistic instrument performance. New virtual instruments can also be created (for example, a saxaphone that sounds when you blow in one end).

There are two common standard types of synthesizers. They fall into the category of either extended or base devices.

- A base level synthesizer device only supports channels/tracks 13-16. The first three of these channels are used for the main song parts (for example, bass, rhythm, and melody). Channel 16 is used as a percussive track (for example, drums). All MPC systems should support the base level.
- Extended level devices support tracks 1-10. The first 9 are for melodic tracks while the tenth is used for percussion.

Most modern synthesizers allow all 16 tracks to be utilized and it doesn′t matter which tracks are used for which instruments.

6.3.1 General MIDI Standard

When assigning various instruments to each track in a MIDI recording, a patch number is used to specify the instrument or sound to use. To help standardize which instruments should be located on individual patch numbers, the general MIDI specification was developed by the MIDI Manufacturer′s Association (MMA).

6.3.2 MIDI Mapper

The MIDI Mapper, which is configured from the control panel, allows non-standard MIDI devices to have their instrument patch numbers reassigned (mapped) to conform to the general MIDI specifications. Percussion key assignments can also be altered.

6.3.3 MIDI Sequencer

A sequencer system is used to record, edit and playback MIDI messages. The sequencer fundamentally acts like a multi-track tape recorder for MIDI instruments. On a computer system the sequencing functions are run by software applications.

6.3.4 When to use MIDI

MIDI is a great alternative to digital audio in the following circumstances:

- File-size is a major consideration. MIDI files are far smaller than wave data files.
- Digital audio will not perform properly. This is often due to the lack of system resources, such as CPU power, disk speed or available RAM.
- You do not require speech overlay.
- Sound quality may be better than digital audio in some cases. This occurs when you have a high-quality MIDI sound source.
- MIDI can be more interactive. MIDI data can be easily manipulated. Details of a composition can be re-arranged.
- Time scaling can be effected without loss of quality or pitch.

6.3.5 Storage Formats

MIDI data can be stored in three different formats: 0, 1, and 2. Multimedia on the Windows PC can only work with formats 0 and 1. Most sequencers can export to these formats. Type 0 is a single track format and is especially good for CD-ROM because it reduces the number of disc seeks and uses less RAM. Type 1 format is for multiple track storage. Both formats have a .MID file extension.

6.4 Digital Movie Formats

Digital movie files are multimedia files that integrate sounds, music, and voices with computer graphics and animation to present information in an exciting, dynamic way.

Movies are made up of a series of still images played in sequence. Each image is called a frame. The number of frames per second is called frame rate, at which a movie is played or recorded.

The movies you can play on your computer are probably different from what you see in the cinema or on TV. Most movie files you can get from the FTP sites are presented in a small window in your computer screen, and they can only be played several minutes, or several seconds. This is because movie files are huge files that take a lot of disk space. If you have a very powerful computer, you will be able to see the real movies on your screen. Actually, some commercial products that can create and playback good quality movies on your computer are already available in the market. If you don′t want to invest your money on these products until you know what they look like, you can get the product demos from the companies′ FTP sites for free.

6.4.1 What You Need to Play Movie Files

To play movie files on your computer, you need a relatively powerful computer.

Hardware requirements:

- Your microprocessor central processing unit, or CPU, must be a 16-Mhz 386SX or higher. A true 32-bit microprocessor such as the 486 is better because it can process and transfer larger amounts of data quickly.
- Your computer must have at least 4 MB of RAM. Of course, the more memory you have, the better.

- The minimum hard disk size is 30 MB; however 80 to 200 MB hard disk drives are recommended. Slow hard disk access time can degrade multimedia performance. A 3.5-inch high-density (1.44 MB) floppy disk drive is also required.
- A sound card with a pair of external speakers or a set of headphones is required to play digitized sound files in high-quality stereo format.
- A VGA video board capable of at least 16 colors at 640x480 resolution. Most standard video boards and monitors meet this requirement. Support for 256 colors is recommended.

Software requirements:

- Audio device drivers for different audio formats
- A video device driver
- Multimedia playback software, and multimedia players

6.4.2 Movie File Formats

Like other files, you can identify movie files by their file extensions. There are only a few movie file formats you can see from the Internet, which are international standard file formats for multimedia.

6.4.2.1 MPEG

MPEG is a very popular movie file format for PCs. MPEG stands for Moving Pictures Expert Group. The members of this group come from more than 70 companies and institutions worldwide including SONY, Philip, Matsushita and Apple. They meet under the International Standard Organization (ISO) to generate digital video standards for compact discs, cable TV, direct satellite broadcast and high-definition television. MPEG meets about four times a year for roughly a week each time. They have completed the *committee draft* of MPEG phase I that is called MPEG I. MPEG I defines a bit stream for compressed video and audio optimized to fit into a data rate of 1.5 Mbps. MPEG deals with three issues: video, audio, and system (the combination of the two into one stream). MPEG is developing the MPEG-2 Video Standard, which specifies the coded bit stream for high-quality digital video. As a compatible extension, MPEG-2 Video builds on the completed MPEG-1 Video Standard by supporting interlaced video formats and a number of other advanced features. Since MPEG deals with three issues, the file extensions by MPEG standards are a little bit different. The most common file extension is .mpg. You will also see:

- .mp2 - MPEG II audio
- .mps - MPEG system
- .mpa - MPEG audio

6.4.2.2 QuickTime

QuickTime is an ISO standard for digital media. It was originally created by Apple Computer Inc. and used in Macintosh. It brings audio, animation, video, and interactive capabilities to personal computers and consumer devices. QuickTime movies are real movies. This standard is much more mature than the MPEG standard. In December 1993, Apple announced that it had begun demonstrating technology that will make future television and multimedia devices more compelling, interactive, and useful for people. Specifically, Apple demonstrated the integration of MPEG technology into applications using QuickTime technology. QuickTime for Windows is available for customers who

use Microsoft's Windows/Dos operating system. QuickTime movies have file extension .qt and .mov. You can play the .mov files on both MACs and PCs.

6.4.2.3 Other Multimedia Video Formats

There are other multimedia file formats. For example, AVI is a video format for Microsoft Windows, and .awa/.awm are Gold Disk Animation. More and more .avi files are available on the Internet. If you have Windows in your computer, you can use Media Player to play (.avi) files. Media Player is in the Windows' accessories group.

6.4.3 Movie Players

To play a movie on your computer, you need a piece of software called the multimedia player, specifically, MPEG player or QuickTime player. These players are also called decoders because they decode the MPEG or QuickTime compressed codes. Some software allow you to both encode and decode multimedia files (for example, to make and play the files. Some software only allow you to play back multimedia files. You have to be very careful to find the correct movie player when you get on the Information Highway. This is because different computers or operating systems use different movie players. There are more movie players for X-Windows and Macintosh machines than for PCs. You run your movie player on your computer and open the movie file within the movie player. Movies on floppy disks should be copied to your hard disk before you play them.

6.5 Multimedia Applications on the Internet

The following area covers some selected multimedia applications that are available on the Internet.

6.5.1 IBM Internet Connection Phone

IBM Internet Connection Phone is the first step in the recent evolution and integration of IBM technologies. IBM-based Internet Connection Phone on real-time voice transfers technology thereby enabling voice transmission through what used to be data-only networks. But IBM technology goes beyond only providing the voice transmission. An IBM research team addressed many of the transmission problems typical of sending voice over data networks. Other incarnations of voice transfer technology have problems with echos and lost packets that lead to transmissions with lots of break up. IBM modified the GSM compression/decompression (codec) algorithm (the European cellular telephony standard) in such a way as to suppress echos and to better control the loss of packets. The new algorithm compresses 8-KHz 16-bit voice samples to 9400 bits per second (bps) leading to clear, near echoless conversations.

IBM researchers continue to integrate other standard codecs such as G.723 and wide-band coders into the improved framework as they become available. The goal is to support a full H.323 network videoconferencing standard. Internet Connection Phone takes full advantage of IBM's MWave technology, the technology that more efficiently processes multimedia and audio data, whenever it can. A computer that has an MWave installed can offload the Internet Connection Phone's compute intensive compression and decompression. This way the computer can do other tasks more effectively while Internet Connection Phone is working. In addition to IBM's innovative technology, IBM is leading the

charge to standardize Internet phones so that users can talk to any Internet phone users independent of the vendor.



Figure 148. IBM Internet Connection Phone. The interface looks like a normal phone device, providing easy operation.

Adding Internet technologies to a company's existing computer network yields an intranet. This intranet has all of the capabilities and features of the Internet but with one major difference: the company has complete control over its intranet. In this case, control means the ability to determine the number of nodes data will pass through when going from point A to point B. It also means the company can base decisions about their network on known information, such as the size of the company, estimated levels of network traffic and acceptable response times.

With the control that an intranet offers, companies can harness the power of Internet technologies to give themselves more function and greater quality of service. In fact, they can virtually guarantee the quality of service. And as Internet technologies advance, companies will have even more power to leverage. The faster response times of an intranet make full function,

multiple-party video conferencing a near-term possibility. Furthermore, intranets put video streaming applications such as viewing live action or long playing videos well within reach. At last, an Internet product that lets you talk, send data and work collaboratively all on a single telephone line.

IBM′s easy-to-use Internet Connection Phone is the first Internet phone product to use ground-breaking IBM technology, therefore providing high-quality voice transmission. In fact, the quality is comparable to the best cellular systems available today. Not only does IBM Internet Connection Phone let you call long distance, with clear, full-duplex speaker phone ability for the cost of connecting to the Internet, but it also can save you time. For example, you can eliminate the time you spend looking up phone numbers and dialing. With Internet Connection Phone, all you do is click on the name of the person you want to call and it connects you. And in the future you will not have to worry if the person has Internet Connection Phone installed, since you will be able to call regular phones and other Internet phone products as we proliferate the telecommunications infrastructure.

You can also easily set up Internet Connection Phone to maximize your efficiency. Internet Connection Phone provides a choice of search algorithms to use on a database that you can customize to meet your needs. For instance, you can organize your private address book by location, relationship or any other criteria. More technical users can go even further by integrating Internet Connection Phone into other applications using the Internet Connection Phone API. If you want to make it easy for people to call you, Internet Connection Phone is your product. People can call you via the Internet by clicking on a link that you set up on your home page. So if a person with an Internet phone can get to your home page, they can get you on the phone.

Internet Connection Phone is easy to use even for people who have never used the Internet. The layout and the help screens provide intuitive guidance on how to accomplish various tasks such as call selection, automatic dialing, and muting. Internet Connection Phone has the major functions we expect from modern telephones and even more. Some of the more advanced features are:

- Call back previous callers
- Configure the phone for computer speakers or headphones
- Select from various servers to find other Internet Connection Phone users
- Adjust microphone sensitivity
- Adjust voice quality

There are other phone technologies available on the Internet, such as:

- WebPhone from NetSpeak
- Internet Phone from Vocaltec
- Televox from Voxware
- Cooltalk from Netscape
- WebTalk from Quaterdeck
- NetPhone from Eletric Magic

6.5.2 Audio on Demand

It is now possible to deliver audio in real-time, on demand, and over the World Wide Web. Indeed it is not only possible; with the advent of faster connections and greater modem speeds, it has become easy. There is a profusion of audio streaming technologies available, such as:

- RealAudio
- Internet Wave
- TrueSpeech
- ToolVox
- AudioLink
- MPEG/CD
- Streamworks
- VDO
- LiveMedia

RealAudio still stands head and shoulders above the others in terms of availability and use but is not an obviously superior product in sound quality and speed. It is the only audio-on-demand software that is currently shipped with Netscape Navigator as a plug-in, and Progressive Networks (developers of RealAudio) have announced a collaboration with Microsoft.

However, VDOLive and ToolVox are also available as plug-ins and other streaming products are likely to follow. It is by no means certain which of the current crop is going to end up as a standard or, indeed, if there is going to be one. As it becomes easier to download software interactively from the Web, there may be less of a need for any one standard to emerge.

6.5.3 Video Conference

Video is a sequence of still images. When presented at a high enough rate, the sequence of images (frames) gives the illusion of fluid motion. For instance, in the United States, movies are presented at 24 frames per second (fps) and television is presented at 30 fps. Desktop videoconferencing uses video as an input. This video may come from a camera, VCR, or other video device. An analog video signal must be encoded in the digital form so that it can be manipulated by a computer.

To understand digital encoding, it helps to understand some background information about analog video, including basic color theory and analog encoding formats. Analog video is digitized so that it may be manipulated by a computer. Each frame of video becomes a two-dimensional array of pixels. A complete color image is composed of three image frames, one for each color component. Uncompressed images and video are much too large to deal with and compression is needed for storage and transmission. Important metrics of compression are the compression ratio and bits per pixel (the number of bits required to represent one pixel in the image). Video compression is typically *lossy*, meaning some of the information is lost during the compression step. This is acceptable though, because encoding algorithms are designed to discard information that is not perceptible to humans or information that is redundant.

Some video-conference technologies available to use on the Internet include:

- Network Video is an Internet video-conferencing tool developed at Xerox/PARC. It is the most commonly used video tool on the Internet MBone. The native nv encoding technique utilizes spatial (intraframe) and temporal (interframe) compression. The first step of the nv algorithm compares the current frame to the previous frame and marks the areas that have changed significantly. Each area that has changed is compressed using transform encoding.

 Either a DCT or a Haar wavelet transform is used. The nv encoder dynamically selects which transform is used based on whether network bandwidth (use DCT) or local computation (use Haar) is limiting the performance. The DCT is desired since it almost doubles the compression ratio. The output of the transform is quantized and run-length encoded. Periodically, unchanged parts of the image are sent at higher resolution, which is achieved by eliminating the quantization step. Typically, nv can achieve compression ratios of 20:1 or more.

- CU-SeeMe is an Internet video-conferencing tool developed at Cornell University. It utilizes spatial (intraframe) and temporal (interframe) compression, with a few twists to optimize performance on a Macintosh, its original platform. CU-SeeMe represents video input in 16 shades of grey using 4 bits per pixel. The image is divided into 8x8 blocks of pixels for analysis. New frames are compared to previous frames, and if a block has changed significantly it is retransmitted. Blocks are also retransmitted on a periodic basis to account for losses that may have occurred in the network.



Figure 149. Video-conference Screen Shots Using Cu-SeeMe (Cornell University)

Transmitted data is compressed by a lossless algorithm developed at Cornell that exploits spatial redundancy in the vertical direction. The compressed size is about 60% of the original (a 1.7:1 compression ratio). The CU-SeeMe encoding algorithm was designed to run efficiently on a Macintosh computer, and operates on rows of 8 4-bit pixels as 32-bit words, which works well in 680x0 assembly code. The default transmitting bandwidth setting for CU-SeeMe is 80 kbps. This number is automatically

adjusted on the basis of packet-loss reports returned by each person receiving the video. About 100 kbps is required for fluid motion in a typical *talking heads* scenerio.

- Indeo is a video compression technique designed by Intel. It evolved from DVI (Digital Video Interactive) technology. Indeo starts off with YUV input, with U and V subsampled 4:1 both horizontally and vertically. Indeo supports motion estimation, using the previous frame to predict values for the current frame and only transmitting data if the difference is significant. Transform encoding is done using an 8x8 Fast Slant Transform (FST) in which all operations are either shifts or adds (no multiplies). Quantization and run-length/entropy encoding are used as in previous algorithms. Indeo specifies that the encoded bit stream be a maximum of 60% of the input data, therefore compression is guaranteed to be at worst 1.7:1.

6.5.3.1 Desktop Video-Conferencing Systems

There are three major platforms for desktop video-conferencing products: Intel-based personal computers running Microsoft Windows or IBM OS/2, Apple Macintosh computers, and UNIX-based workstations running the X Window System. Unfortunately, there is currently very little interoperability among products and platforms. Products are evolving towards conformance to the emerging desktop video-conferencing interoperability standards. All systems require hardware that captures and digitizes the audio and video. Video is typically input in NTSC or PAL formats.

Most systems have some sort of graphical user interface that assists in making connections to other parties, usually utilizing the paradigm of placing a telephone call. Many products allow you to store information about other parties in a phone book or Rolodex format. Systems commonly have controls to adjust audio volume, picture contrast, etc. Many systems have controls that allow you to adjust the transmitted bandwidth for video to minimize traffic on a network. An additional feature found in most systems is a shared drawing area usually called a whiteboard which is analogous to the whiteboards found in many conference rooms and classrooms. These whiteboards commonly allow participants to import other graphics such as images and to make annotations. Whiteboards are good for simple sketches, but fine detail is difficult to achieve using a mouse.

Many systems allow an easy way to transfer files between participants. Some systems allow application sharing, which enables a participant to take control of an application running on another participant's computer. The usefulness of application sharing is often demonstrated with an example of sharing a spreadsheet or word processor program to facilitate group collaboration.

6.6 Multimedia Glossary

8-bit sound: Sound which is approximately equal in quality to broadcast radio sound. (See Sample size for further explanation.)

16-bit sound: Sound which is approximately equal in quality to standard audio Cds. (See Sample size for further explanation.)

Access time: The time it takes for the computer to begin reading from or writing to a storage device such as a hard drive or CD-ROM drive.

ADC: Analog Digital Conversion - The method of converting analog data to digital data (as in analog-to-digital sound).

ADPCM: Adaptive Delta Pulse Code Modulation - A way of storing analog sound in a compressed digital format.

AGC: Automatic Gain Control - A process that levels out high and low levels of sound to improve the consistency of the recording.

Audio board: An expansion board that you put inside a PC to improve the quality of the PC′s sound output. Also called a sound board or sound card.

AVI: Audio Video Interleave - A specification that allows for the capture and storage of video and waveform audio in a single data stream. Because of speed and memory limitations, AVI offers only rough animation, not full-motion video.

CD or Compact Disc: An optical read-only disc that is used to store digital audio, data, or video. CD-ROMs provide about 600 MB of storage space.

CD-audio: Sounds that have been digitized at a sampling rate almost high enough to duplicate reality. CD-audio is the same format and quality as the discs you play on your CD player.

CD-DA: Compact Disc - Digital Audio - CD-quality audio that comes directly from a CD-ROM or an audio CD.

CD-I: CD Interactive - An interactive audio/video computer system developed by Sony and Philips.

CD-ROM: A type of compact disc that stores digital data.

CD-ROM drive: A device that reads from and writes to CD-ROMs. A CD-ROM drive can be installed in the computer (internal drive), or it can be connected to the computer (external drive.) A CD-ROM drive lets you store data or play sound directly from the drive.

CD-ROM XA: CD-ROM Extended Architecture - An extension of the CD-ROM standard that permits sound and video data to be combined and synchronized with animation.

Composite video: A color video signal that contains all of the color information in one signal.

Compression: A process that allows data to be stored or transmitted using less than the normal number of bits.

DAC: Digital Analog Conversion - The method of converting digital data to analog data as in analog sound to digital sound.

DSP: Digital Signal Processor - A processor that can be programmed to perform certain tasks such as compression or sound effects.

Digital audio: Data that is recorded and processed to create sound for editing and playback.

Digital video: Video that is recorded and processed for editing and playback.

Digitize: The process of converting analog data to digital data.

Dual Speed: A CD-ROM drive that accesses data at 300 kbps. This is twice as fast as a standard audio CD player or single speed CD-ROM drive.

DVI: Digital Video Interactive - A form of video compression from Intel.

Dynamic range: The span of volume between the loudest and softest sounds in an audio recording. Sample size affects dynamic range. 16-bit audio yields a dynamic range of 96 dB, and 8-bit audio yields 48 dB.

External CD-ROM drive: A CD-ROM drive that is installed outside the computer and is connected by a cable to the computer.

Filtering: A digital conversion process that improves the fidelity of audio recording.

FM synthesis: A technique for synthesizing sound that uses a combination of modulated sine waves to produce different wave forms.

Full-motion video: Video reproduction at 30 frames per second for NTSC signals or 25 frames per second for PAL signals.

Full-motion video board: An expansion board that you put inside a PC that allows you to capture, digitize, and compress multiple frames from an NTSC video source. The frames can be stored on a hard disk or other storage device.

Interframe compression: A form of video compression that compresses full-motion video by analyzing each frame of a video, determining which frames duplicate previous frames, and deleting the duplicates.

Internal CD-ROM drive: A CD-ROM drive that is installed inside the computer.

Intraframe compression: A form of video compression that compresses full-motion video on a frame-by-frame basis.

JPEG: Joint Photographic Experts Group - A form of intraframe compression that offers a maximum compression ratio of 20 to 1.

LMSI: A proprietary interface developed and used by Philips to connect Philips CD-ROM drives to a PC.

Lossless compression: A type of data compression that makes it possible to recover the original data with no loss of image quality.

Lossy compression: A type of data compression that sacrifices some of the original data in return for higher compression ratios than can be achieved with lossless compression.

MCA: Media Control Architecture - A specification developed for addressing various multimedia devices from Macintosh computers.

MCI: Media Control Interface - A platform-independent multimedia specification that provides a consistent method for controlling multimedia devices.

.MID: MIDI file extension.

MIDI: Musical Instrument Digital Interface - A digital communications standard that lets electronic musical instruments and computers communicate with each other. MIDI files are typically saved with a .MID file extension.

MPEG: Motion Pictures Experts Group - A form of interframe compression.

MPU-401: A standard MIDI interface that features its own CPU for processing some MIDI data without taxing the main computer's resources.

Multimedia: The use of two or more media types (motion video, audio, still images, graphics, animation, text, etc.) to communicate information.

Multimedia extensions: Tools in Windows that enable developers to create multimedia applications.

Multimedia PC (MPC): A standard computer configuration recommended for multimedia.

Multimedia upgrade kit: A complete package of hardware (CD-ROM drive, sound board, and speakers) and software that adds multimedia capabilities to your PC.

NTSC: National Television Standards Committee - The standard broadcast signal received by TV in the United States.

PAL: Phase Alternation Line - The standard broadcast signal received by TV in many European countries.

PCM: Pulse Code Modulation - A digitization technique that places audio on a tape.

Photo CD: A product developed by Eastman Kodak that places photos on a compact disc and allows users to view them on their television or computer.

Photo CD compatibility: A product that displays photos stored on a compact disc. Photo CD-compatible products can support what is described as single session (capable of displaying only one set of photos from the CD) or multisession (capable of displaying more than one set of photos from the CD).

RAM: Random Access Memory - The part of a computer's memory used to write data to and read from a disk. When you work on a computer, the information displayed on the monitor screen is stored in RAM.

RTV: Real Time Video - A form of interframe compression that allows for compression rates of up to 150 to 1.

Sample size: The number of bits used to store the recorded sound's amplitudes. It is also referred to as resolution. The sample size is measured in bits and governs the difference in volume between the softest sound and the loudest sound that can be recorded and played back. The sample size of standard audio CDs is 16 bits, and the sample size for standard broadcast radio is 8 bits. 16-bit audio allows 65,536 loudness levels, whereas, 8-bit audio allows 256 loudness levels. Combined with sample rate, it provides a measure of how closely a sound that is recorded and played back will match the original sound source.

Sampling rate: A measure of how often sound is converted from an analog waveform to numbers. The sampling rate is measured in samples per second and governs the highest and lowest frequencies of sound that can be recorded and reproduced. Standard audio CDs use a sampling rate of 44 kHz. The 44.1 kHz sampling rate captures 44,100 (amplitude samples) picture of sound per second. Combined with sample size, sampling rate provides a measure of how closely a sound that is recorded and played back will match the original sound source.

SCSI: Small Computer System Interface - An industry-standard connection for hardware devices.

Signal-to-noise ratio: The ratio of the desired signal (for example, music) to extraneous noise (such as background hiss), expressed in decibels.

Single speed: A CD-ROM drive that accesses data at a speed of 150 kbps. This is the speed at which standard audio CDs can be read. Single speed is the standard speed for CD-ROM drives.

SLCD: A proprietary interface developed and used by Sony to connect Sony CD-ROM drives to a PC.

S-Video: A type of video signal that transfers light and color separately, using multiple wires. S-video delivers a higher quality picture than formats such as NTSC which encodes the data.

Synthesized audio: Audio output from a synthesizer.

Synthesizer: An electronic musical device that generates sound.

Transfer rate: The time required for data to be transferred from the hard drive (or CD-ROM drive) to the computer's CPU.

Triple speed: A CD-ROM drive that accesses data at 450 kbps. This is three times as fast as a standard audio CD player or single-speed CD-ROM drive.

Video capture board: An expansion board you put inside a PC that allows you to capture a single frame from an NTSC source and save it on your hard disk.

Video for Windows: A software program that lets users play video on their PC without additional hardware.

Video pass through: A television or other video source connected to the computer to play video on the computer screen.

.WAV: The file format for waveform audio.

Waveform audio: A form of digital audio that is stored in a format the PC can understand and manipulate. Waveform audio is generally stored with a .WAV extension.

Chapter 7. Existing Gateways

In this chapter the discussion of gateways describes the application interfaces that enable WWW applications to access data stored in your local databases. It is not always easy to create this interface to fit your specific needs. However, there are some solutions that have already been developed to aid in the implementation of this application interface.

Although solutions or programs exist to fit specific needs, another facet of this development is the support and flexibility required by the owners and users of the environment. The development of these solutions has prompted some companies to identify the problems and then create solutions that avoid the pitfalls. The results are true interfaces from the WWW to existing environments such as database systems, mainframe applications, and other specific environments.

Some application interfaces, such as the DB2 gateway, have been created utilizing the extensive database knowledge to development better interfaces between the Web and database systems.

7.1 DB2WWW Gateway

The database is the main component in most systems. DB/2 is the most important database because of the amount of data involved.

IBM has developed tools to make data access easier from the WWW or an intranet, making browsers a powerful database front-end for applications.

The DB2 Internet gateway allows programmers to create applications with a simple tool, and without the expensive programming time that a database system program requires.

DB2 gateway works by interacting with the server and the database system, at the same time the server makes its own interaction with the Web browser, as shown in Figure 150 on page 300. DB2 can interact directly on the database or it can use the DB2 Software Development Kit to access the database systems; this approach allows you to put information that is not in your Web server. For example, an AS/400 database can be accessed from a Web server that is on an AIX machine.

The DB2 gateway is available for the following platforms:

- OS/2
- AIX
- OS/400
- Windows NT
- Sun Solaris
- HP-UX
- MVS

© Copyright IBM Corp. 1996

It is available for several languages such as English, simplified Chinese, traditional Chinese, German, Spanish, Italian, Japanese, Korean and Portuguese-Brazilian in the OS/2 platform.

To download the DB2WWW on the current version go to `http://service.boulder.ibm.com/pbin-usa-demos/getobj.pl?/demos-pdocs/wwwdb2dnld1.html`. To get information about the features, installation processes, etc., go to `http://www.software.hosting.ibm.com/data/db2/db2wgafs.html`.

To understand how the macro files on DB2 works you must know the HTML specifications and SQL.

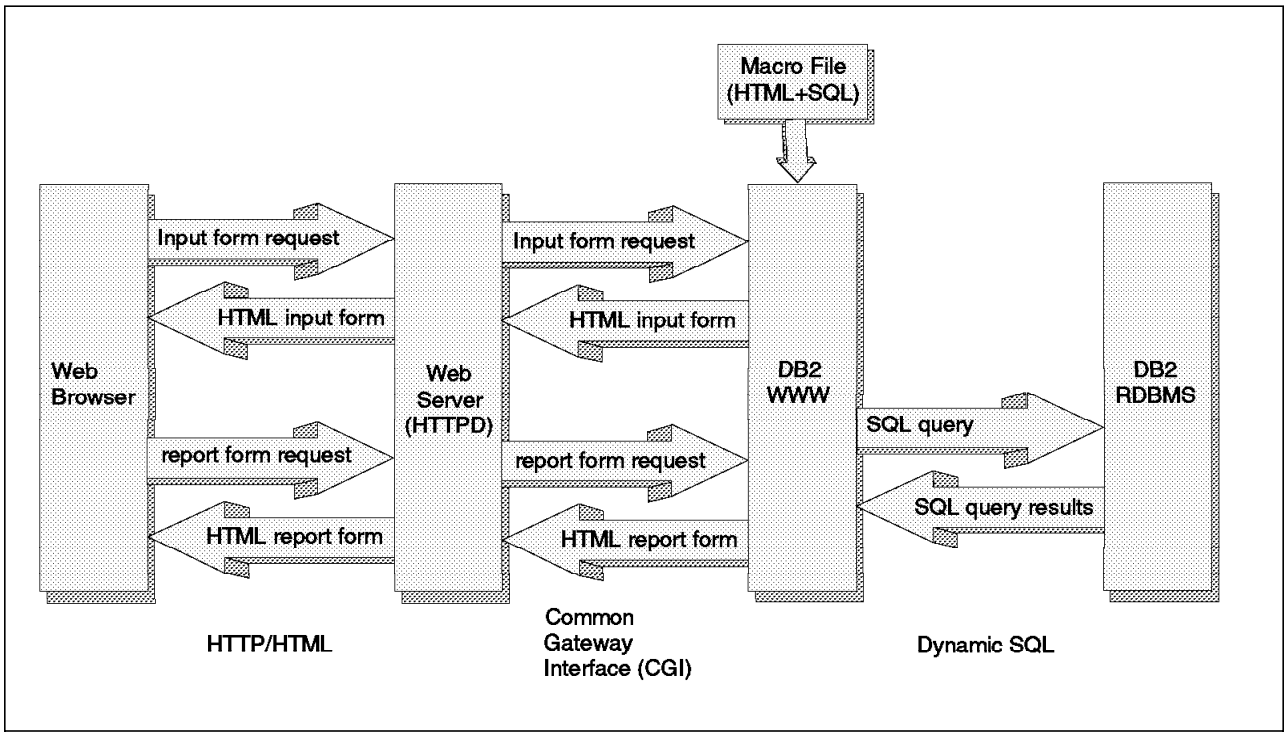


Figure 150. The DB2 Data Flow When Used by the DB2WWW Gateway

7.1.1 Installation

The procedure is variable depending on the platform. Most of the time the program comes in a compacted file (for example, .zip for PC or .tar.Z for UNIX systems). Once you decompress your files, you must check the following:

1. The DB2WWW executable program must be in the cgi-bin directory, or equivalent.
2. There must be a db2sql.bnd file in the same directory as the executable.
3. The DB2WWW.ini has to be in the home directory.

The installation depends on the platform. For example, AIX has the SMIT tool that allows the DB2 gateway to be installed. OS/2 DB2WWW has an install program that appears when the .zip file is decompressed. If you change your home page location the DB2 gateway installation program will not seek for the new location on the http configuration file.

7.1.2 Configuring DB2WWW

Before you start the central part of this section (how to use the DB2WWW gateway) you have to configure it to get it to work. There are two special files to configure the database gateway: the initialization file and the bind file.

The initialization file has two lines, both of which are paths to specific places: the macro library directory and the bind file.

The format of the initialization file is:

```
MACRO_PATH C:\DB2WWW\MACRO
BINDFILE C:\WWW\CGI-BIN\db2sql.bnd
```

The different operating systems can accept or refuse those kind of paths, so they have to be changed to the correct format depending on the operating system. The example above can work with either NT or OS/2.

A bind file is a file used by DB2 in order to find a better way to access the database. This file must be updated for every new database you want to access, and can have unlimited paths for the databases. To update the bind file you must follow these steps:

1. If you are in a UNIX system, log on as a user who can have access to the database (usually the instance owner).
2. If you are using OS/2, you may access your database access program and use the DB2 command line to generate the bind.
3. Use the `bind` command in order to add the new database to the specified file.
4. Log out from the database and use the new bind file in the db2www.ini file.

7.1.3 The Macro File

Once your DB/2 gateway has been configured, you can begin to work on your applications, and you are ready to write your DB2WWW macros.

A macro has four different sections:

1. Define section
2. HTML input section
3. SQL section (could be one or more SQL sections)
4. HTML report section

The macro files are plain text (ASCII) files. These files contain a special syntax in order to get all of the variables and imbed them into an HTML response from the gateway. Every section begins with the symbols `%NAME_OF_SECTION{` and ends with the `%}` symbols. The comments should be in a separate section without a name.

7.1.3.1 Define Section

This section contains all of the variable definitions. The most important variables are those relative to the database, including:

Table 26. Variables of DB2WWW	
Variable Name	**Meaning**
DATABASE	The name of the database to be accessed; it has to be included on every macro.
PASSWORD	The correct password to access the database; it is associated with the LOGIN variable.
LOGIN	Indicates the user ID for the database.
SHOWSQL	It contains a flag to show the SQL command. The default value is set to *no*. To display the command set this variable to *yes*.
DB_CASE	*UPPER* or *LOWER* are used to convert all letters to upper or lower case. The default value is null (none of the two conversions occur).
RPT_MAX_ROWS	Maximum number of rows displayed by the browser. Values such as *0, all* and *ALL* can be set.
ALIGN	Leading or trailing spaces are used to create a table so the values are aligned properly in the query results.

You can use a block to define a variable with a value longer than one line. For example:

```
%DEFINE{
    DATABASE="CostumerDB"
    LOGIN="MyUSSSSERID"
    PASSWORD="Password"
    Variable1={This is a multiple line
                Sting on a DB2WWW macro File%}
%}
```

You can also use %DEFINE to declare only one variable. For example:

```
%DEFINE DATABASE="celdial"
```

7.1.3.2 HTML Input Section

To invoke the DB2 gateway, you must use the following link:

```
http://.../cgi-bin/nameofdb2www/command
```

where nameofdb2www can be db2www in the UNIX platform or db2www.exe in the PC's operating systems such as OS/2 or Windows NT. The command can be report or input options. If input is chosen, the %EXEC_SQL commands are not executed and the %HTML_REPORT is not displayed. If the report option is chosen, the %INPUT_HTML section is not displayed. You must put the form in both if you want to get the feedback, such as in a search engine.

```
%HTML_INPUT{
<TITLE>DB2 WWW Company Information Query</TITLE>
<img src="/icons/head1.gif">
<P ALIGN=center>
<A HREF="/saledoc.htm">
About this macro...<IMG SRC="/icons/bigqmboo.gif" ALIGN=middle>
</A>
<br><P>
<H1>Welcome to the Celdial database</H1>
<p>
This query retrieves information about a company, including the
company name.
<hr>
Choose which additional fields you would like to see in the results:
<FORM METHOD="POST"
ACTION="/cgi-bin/db2www.exe/saleqadd.d2w/report">
<INPUT TYPE="hidden" NAME="field" VALUE="$(tc).custname">
<p>
<SELECT NAME="field" MULTIPLE SIZE=4>
<OPTION VALUE="$(tc).contact,$(tc).con_phone">Contact Name and phone
<OPTION VALUE="$(tc).con_addr">Contact address
<OPTION VALUE="$(tc).con_country">Country
<OPTION VALUE="$(tc).custno">Customer number
</SELECT>
<hr>
Enter the company name and the contact name in the input fields
provided below.  You do not need to enter all of the characters of a
name.  For example, you can use "Mer" instead of "Meridien".
<p>
<pre>
Company Name: <INPUT TYPE="text" NAME="INPUT_CUST_NAME" VALUE="Meridien" SIZE=25>
<br>
(Examples: Meridien Elec, Royal Hardware, Holmes, Holiday, Hollister)
<p>
Contact Name: <INPUT TYPE="text" NAME="INPUT_CONTACT_NAME"
VALUE="Alfredo Bay" SIZE=15>
<br>
(Examples: Alfredo Bayon, Arnie Podel, Zoltan, William, Yutaka)
</pre>
<hr>
Select which type of query you wish to perform using the company name
and contact name above:
<p>
<INPUT TYPE="radio" NAME="INPUT_ANDOR" VALUE="AND" CHECKED> List all
companies using <strong>both</strong> company name and contact name
(logical <strong>and</strong>)<br>
<INPUT TYPE="radio" NAME="INPUT_ANDOR" VALUE="OR"> List all companies
using <strong>either</strong> company name or contact name (logical
<strong>or</strong>)
<hr>
```

Figure 151 (Part 1 of 2). Input Section for the saleqadd.d2w File

```
Show SQL statement on output? <INPUT TYPE="radio" NAME="SHOWSQL"
VALUE="YES"> Yes
<INPUT TYPE="radio" NAME="SHOWSQL" VALUE="" CHECKED> No
<p>
<INPUT TYPE="submit" VALUE="SUBMIT QUERY"> <INPUT TYPE="reset"
VALUE="Reset">
</FORM>
<p>
<hr>
<p>
Other pages of interest:
<P>
<A href="/celdemo.htm">DB2 WWW Connection Demonstrations</A>
<br>
<A href="http://www.software.ibm.com/data/db2/db2wfac2.html">
DB2 WWW Connection Home Page</A>
<br>
<a href="/celdial.htm">DB2 WWW Connection Celdial Demonstration</A>
<P>
<hr>
<b>
[
<a href="http://www.ibm.com/">IBM home page</a> |
<a href="http://www.ibm.com/Orders/">Order</a> |
<a href="http://www.austin.ibm.com/search/">Search</a> |
<a href="http://www.ibm.com/Assist/">Contact IBM</a> |
<a href="http://www.ibm.com/Finding/">Help</a> |
<a href="http://www.ibm.com/copyright.html">(C)</a> |
<a href="http://www.ibm.com/trademarks.html">(TM)</a>
]
</b>
%}
```

Figure 151 (Part 2 of 2). Input Section for the saleqadd.d2w File

Figure 151 on page 303 shows the use of the input section to create forms. The screen shown for the browser should be as shown in Figure 152 on page 305.

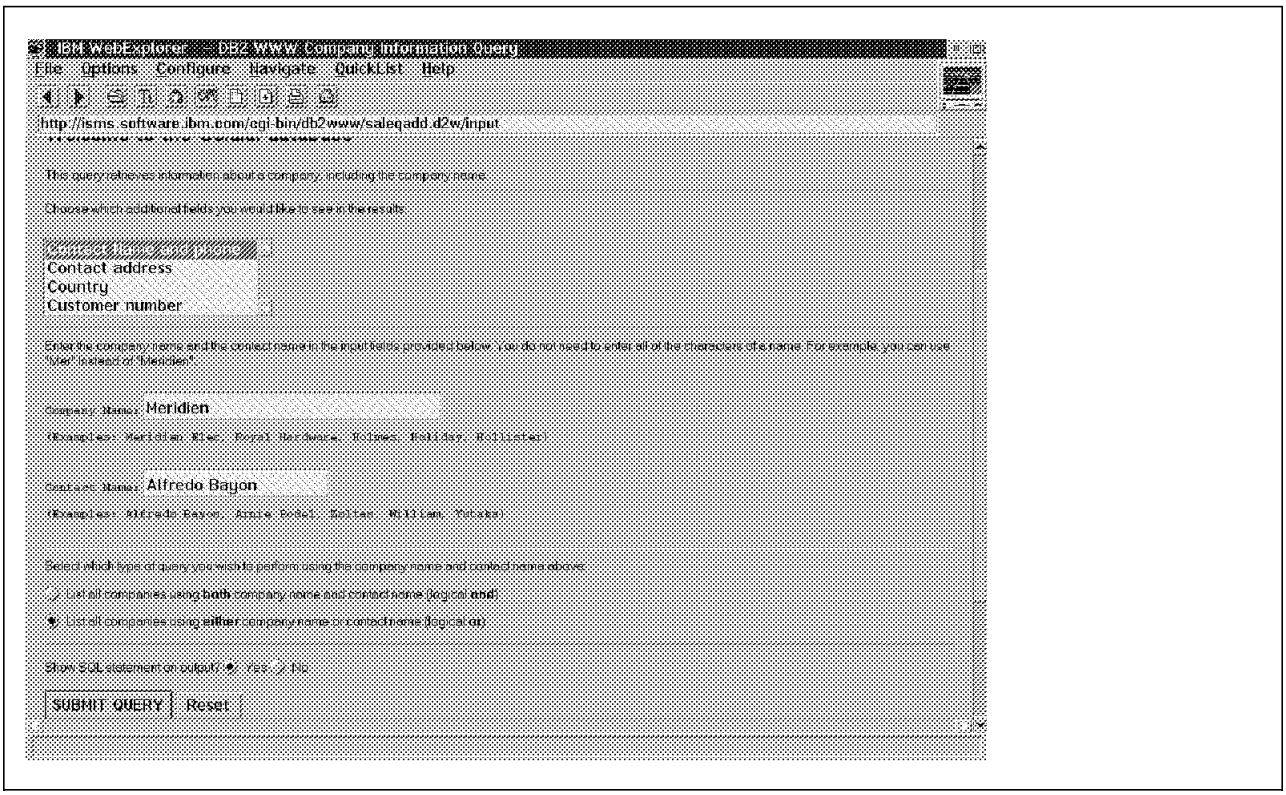


Figure 152. Form of the Input Section

7.1.3.3 The SQL Section

SQL is the most powerful tool to create queries and update databases. The commands received by the Database Management System (DBMS) are processed and sometimes are passed to another system that uses a different database format. Heterogeneous DBMSs are used in a wide range of enterprises, and the common language they use is SQL.

In the %SQL section you must enter one SQL statement and the format you are going to use to display the data.

The %SQL_REPORT and %SQL_MESSAGE are two subsections. The first one allows you to control the data returned by the database system, since you can or cannot be using DB2, if the return code indicates no error or warning. The second one allows to change the messages in case of error or warnings appearance. The format for the entire %SQL section is:

```
%SQL (sql-section-name){
    Any SQL
    on multiple lines.

    %SQL_REPORT{
    Any valid header HTML or column variable names
    returned from the query.
    %ROW{
    Any valid HTML with special variables
    to display once for each row returned.
    %}
    Any valid HTML footer HTML.
    %}
```

```
%SQL_MESSAGE{
+SQLCODE: "warning message" : exit or continue
+SQLCODE: "warning message" : exit or continue
-SQLCODE: "error message"
-SQLCODE: "error message"
default: "default message"
%}
%}
```

The SQL in a section is executed when it is called by %EXEC_SQL in the HTML report section.

If an error or warning occurs in an SQL command, the execution terminates and a return code is given.

You must decide if you want the application to continue after receiving a warning message from an SQL command. Information dealing with these issues is in SQL Message Subsection. This example returns a list of all products in PRODTABLE and orders them using a variable specified through an HTML form in the HTML input section:

```
%SQL(prodList){
   SELECT MODNO, MANUF, COST FROM PRODTABLE
   ORDER BY $(ordby)
   %}
```

Note: DB2 for OS/2 Version 1.2 and DB2/6000 Version 1.2 do not support SQL containing tabs or carriage returns.

The SQL Report Subsection: This subsection gives you the ability to customize the query output using HTML formatting. If you have no SQL report subsection, a default table is displayed with column names at the top.

All text and graphics before the %ROW declaration is header information and is displayed before any information from the SQL query. Following the SQL query processing, the column names are placed in special variables N*i*, N_*column-name*, and NLIST.

The ROW subsection contains information displayed once for each row returned by the SQL query.

Information, including text and graphics, following the ROW subsection is footer information and is displayed once after all rows are displayed.

This are some variables that can help you to create your DB applications with the DB2 gateway.

Table 27 (Page 1 of 2). Variables used for the DB2 gateway that cannot be changed.

Variable	Meaning
N1, ..., Ni	The name of the columns in the report. These variables are only valid within the SQL report section.
V1, ..., Vi	The values for each field of a row returned by an SQL query. They are only valid inside the ROW section. The values change as each row is retrieved.

Table 27 (Page 2 of 2). Variables used for the DB2 gateway that cannot be changed.

Variable	Meaning	
N_column-name	The name of the specified column name. If the column name does not exist, this variable is not defined. For example, the value of $(N_ZIP) is ZIP.	
V_column-name	The value for the specified column name for the current row. This variable is not defined if the column name does not exist. For example, the value of $(V_ZIP) might be 98109.	
NLIST	This is a special list variable that contains all the column names from the result table. The default separator is a space, but you can specify another separator in the DEFINE section this way: %DEFINE %LIST ″	″ NLIST A query returning names and phone numbers might have $(NLIST) with this string: LAST FIRST AREA NUMBER. This variable is most helpful when creating tables in HTML 3.0.
VLIST	The field values for each row of the result table. The default separator for the names is a space, but you can specify another separator in the DEFINE section. For example: %DEFINE %LIST ″	″ VLIST. A query returning names and phone numbers might have $(VLIST) for the first row with this value: ANH TERESA 408 555 9876 This is most useful when creating tables in HTML 3.0.
ROW_NUM	The current number of rows retrieved from the query. When the last row is returned, this variable contains the total number of rows returned.	
NUM_COLUMNS	The number of columns returned by the SQL query.	
SQL_CODE	Contains the SQL warning or error from the SQL query. Successful SQL queries result in 0.	

SQL Message Subsection: This subsection allows you to customize error and warning messages from SQL commands. If you place this declaration inside an SQL section, it is local only to the SQL command in that section. If it is outside of all SQL sections, it is global to the entire macro.

Create a table of SQL codes and specify the information to display following each SQL code. The default error message is shown when an SQL code not in the declaration is returned by the special variable SQL_CODE. For positive SQL codes, you have the option of exiting or continuing. Table 28 shows how different conditions are handled:

Table 28 (Page 1 of 2). Results Following an SQL Warning or Error

SQL Return Code	Local or Global SQL Messages Declaration Exist	No Local or Global SQL Messages Declaration Exist
Positive	Warning displays, procedure continues or stops.	DB2 Message displays, process ends.

Table 28 (Page 2 of 2). Results Following an SQL Warning or Error		
SQL Return Code	**Local or Global SQL Messages Declaration Exist**	**No Local or Global SQL Messages Declaration Exist**
Negative	Error message displays, processing ends.	DB2 default message displays, processing ends

You can have as many SQL sections as you want, and you call them in the HTML_REPORT section.

7.1.3.4 HTML_REPORT Section

This is the part where you are going to create the HTML page based on the results of the Web based on the queries.

This section is where you call the SQL query. The section is executed when DB2 World Wide Web is started in the report mode, often from the HTML input section of the macro.

```
%HTML_REPORT{
    any valid HTML text
    %EXEC_SQL(SQL section name)
    any valid HTML text
    %EXEC_SQL(SQL section name)
    any valid HTML text
    .
    .
    .
    %EXEC_SQL(SQL section name)
    any valid HTML text
    %}
```

You can specify any HTML and include any variables from the DEFINE section in the HTML code. Use input from the HTML form to override variables in the %DEFINE section. When an %EXEC_SQL line is encountered, the SQL section matching the name or defined variable is called. Using a variable for the SQL section name is an easy way to allow customers to select a query to perform.

If you do not specify a section name, all unnamed SQL sections are executed in the order they appear in the macro.

Here is a simple example of what an HTML report section might look like. You can define the variable query in the DEFINE section, or have the application user specify a value in the input section.

```
%HTML_REPORT{
    <HEADER>
    <TITLE>Database query results</TITLE>
    </HEADER>
    <IMG SRC="gifs/logo.gif" ALIGN=MIDDLE>
    <BR>
    %EXEC_SQL ($(query))
    <HR>
    <A HREF="/cgi-bin/db2www.exe/query.d2w/input">
Submit another query</A>
    <br>
    <A HREF="www.celdial.com">Home page</A>
%}
```

A good idea can be to add a form before the results are shown to let the user generate another request from the same screen.

There are some interesting tricks you can find in the documentations that come with the gateway. Hidden variables, for example, can be useful in maintaining security on your pages. Users will not know what the variables will be, even if they browse the HTML file. The variables will be shown as a `$(variable)` variable. These kinds of tricks can help you to develop fine applications on your Web server. Use the hidden variables and the conditional statements for better applications.

7.1.4 Accessing Non-DB2 Databases with DB2WWW

The DB2 gateway allows you to connect to databases different from the DB2 standard using Data Joiner instead of using the DB2 database system.

7.2 Other Database′s Gateways

Most databases have their own gateway. Sometimes the same company develops this tool and sometimes it is created by a third party, but the result is the same.

The flexibility of each tool depends on the approach that every company puts on the products.

Oracle, Sybase and Informix are powerful databases used by corporations to keep data. All of them have different characteristics.

7.2.1 Oracle

Oracle is developing more than a simple database solution. Oracle has the solution for the data management and Web server integration. However, the solution is not available in a wide platform environment. The Oracle Personal Edition is one of the most popular databases available. It is a cross platform (Windows, NetWare, PowerMac, and OS/2) and can be accessed easily from the different C compilers with the included products, making the applications work harder to create CGIs. However, the Workgroup server provides complete Web integration providing the Web server and the tools to create enabled Web applications in an easy way. The product bundles the Oracle Web server making your applications appear in the Internet, enabling store procedures to be invoked by the server to generate dynamic pages.

For more information about Oracle products, available gateways and servers, go to the URL: `http://www.oracle.com`.

7.2.2 Sybase

Sybase is an important database in the business world. It is available only on Digital, Windows NT, HP/UX and Sun platforms. The gateway used by the database to generate the integration with the server is called Web.sql.

Sybase′s gateways provide complete integration with the server′s API, making this gateway an extension of the server to improve the performance. At the time of writing, this gateway was only available for SUN Solaris and Silicon Graphics IRIX. You can download a trial version from the Internet for these platforms. The

gateway is planned to be available for HPUX 9.0 and Windows NT. At the moment you can download an alpha version from the Internet.

Sybase has an agreement with Netscape to use the Netscape Commerce Server as a part of their solutions.

For more information on Sybase products go to the following URL:
`http://www.sybase.com`

7.3 MQSeries Gateway

The MQSeries Internet Gateway provides a bridge between the synchronous World Wide Web and asynchronous MQSeries applications. Interaction with the gateway is via HTML fill-out form POST requests. The form needs to identify the target queue and queue manager names that the application servicing the requests will be using. The MQSeries application receiving the request will need to be able to generate HTML pages to return to the gateway.

7.3.1 Software

The gateway has been tested on the following operating systems and Web servers:

- AIX 3.2.5 with NCSA HTTPD Version 1.4
- AIX 4.1.4 with NCSA HTTPD Version 1.4
- OS/2 Warp with IBM Internet Connection Server 4.0 and VisualAge V3.0

7.3.2 Installation

Installation will depend on the server that is being used. Web servers tend to have a default path for CGI executable binaries and another path for HTML documents. However, this path can also normally be configured to be whatever the installer of the server desires. Hence, the CGI programs and the sample HTML files should be placed in the appropriate directories according to the Web server being used and its configuration.

The following files should be put in the directory for CGI programs:

- MQGate
- MQHost
- timedout.html
- MQQueueB
- MQGate.ini
- amqwput
- amqwget

The OS/2 version of the gateway also has the file cgilib.dll, which should also be placed in the directory for CGI programs.

The following files should be put in the root HTML document directory:

- The Gateway home page, MQGate.html
- The user guide, igp.html

- The host name sample HTML file, MQHost.html
- The queue browser sample HTML file, MQQueueB.html
- The put sample, amqwput.html
- The get sample, amqwget.html

7.3.3 Gateway Components

- MQGate

 This is the CGI program that HTML pages should specify in the form action URL. The program essentially just performs the MQSeries API calls to:

 – Connect to the queue manager specified

 – Open the requested queue

 – Put a message whose data content is the stdin data received by the program on the queue

 – Open the gateway's reply queue

 – Wait for the response message to arrive

 – Write the message data content to stdout

- Web Server

 The Gateway should work with any CGI-capable server. However, it has only currently been tested with NCSA HTTPD 1.4.

- Web Browser

 A form-capable browser.

- Gateway.Reply.Queue

 This is the default queue on which the gateway will be expecting a reply message to any request messages that it has made. The name can be configured by using the MQIGwReplyQueue field in the MQGate.ini file.

 This is the destination queue for the message generated by the MQGate. Each application can have its own queue, or several applications can share the same queue.

- Application

 The application needs to be able to process MQSeries messages that have CGI style separators and delimiters and be able to produce HTML format output messages.

The Web server and Web browser are not supplied. A sample application is provided as well as a script to create the queues it needs.

7.3.4 MQGate

The basic function of this CGI program is to convert the CGI data received on stdin to an MQSeries message, put it on a queue and then wait for a response. The gateway is also the crossover between the synchronous world of the Internet and the normally asynchronous world of MQSeries. This difference is handled by having a user defined time out limit on the MQGET of a response message, when the wait limit is exceeded an HTML page is sent to the client. The default page has two action buttons. One to cancel the wait, the other to perform another get wait.

7.3.4.1 MQSeries Queue Manager and Queue

The Gateway needs a target Queue Manager and Queue. These can be provided in variables passed from the HTML, typically as hidden fields, or as defaults in the gateway configuration. The variables are MQIGwQueueManager and MQIGwQueue. The Gateway will use a variable from the HTML if present, if not then the Gateway will search MQGate.ini for the variable.

The HTML coding to generate the name/value variables would typically be:

```
<input type=hidden name=MQIGwQueueManager value="My.Queue.Manager">
<input type=hidden name=MQIGwQueue value="My.Queue">
```

7.3.5 Configuration

The Gateway needs at least two MQSeries queues, one to receive reply messages on and one on which to put outgoing messages destined for an application. The queue that MQGate is to put to is determined by the HTML form, as described above. The name of the Reply Queue is held in the MQGate.ini file. The default MQGate.ini file sets a Reply Queue name of Gateway.Reply.Queue.

7.3.5.1 Gateway Timeout Form

The default HTML timed out page is timedout.html. This page can be replaced, but the replacement should contain the two submit buttons from the default page.

7.3.5.2 MQGate.ini

This .ini file is used to specify configuration values for the gateway.

MQIGwWaitInterval Used to define the wait limit on an MQGET performed by the gateway. If no file is found then the default wait limit is 30 seconds.

MQIGwReplyQueue Defines the queue on which the gateway will wait for responses and that it will put in the ReplyToQueue in the message descriptor of any request messages.

MQIGwQueue This can be specified to give a default queue to be used by the gateway for a request message, if there was no MQIGwQueue variable passed from the HTML.

MQIGwQueueManager This can be specified to give a default queue manager to be used by the gateway, if there was no MQIGwQueueManager variable passed from the HTML.

7.3.6 Host Name Sample Application

This sample application, which is included with your gateway, shows you how to return the TCP host name of the machine on which it is running. MQHost will sit in a get wait on the queue specified when it is invoked. Once it receives a message on the queue, it will construct a response message that contains the host name. In this trivial sample no checking of the CGI content is done by the application; receiving a message is all that is needed. The data content of the message is of the form:

```
Content-Type: text/plain

this.machines.host.name
```

This message will be put on the reply queue specified by the request message, which is set by configuring the gateway. The gateway will get the message and

write the content to stdout for the Web server to process. The invoking Web browser will then receive the data and display the name as a simple line of text.

The sample consists of the following files:

- MQHost.html

 This is the HTML form to submit the action.

- MQHost

 This is the executable for the application.

7.3.6.1 MQHost

The application is a long-running server that is started by typing:

```
MQHost QueueName QueueManagerName
```

The server can only be terminated by a kill.

7.3.6.2 MQHost.html

This HTML form only has a submit button visible. It also has hidden fields that indicate the target queue and queue manager the gateway will be using to put messages on for the application. The default setting of these is:

Queue MQHost.Queue

Queue manager Set to blank for the default queue manager

To use the sample, either create a queue of this name and use this and the default queue manager name to invoke MQHost, or edit the HTML to use any other queue and queue manager names required.

7.3.7 Queue Browser Sample Application

This sample application provides simple remote queue browser capability. The sample application needs to be running on the queue manager where the queue to be browsed is, but this can be a different queue manager or system from that where the Web server is running.

The MQQueueB sits in a get wait on the queue specified when invoked. Once the MQQueueB receives a message on this queue it needs to decide what to do with it. There are two basic messages that it can receive: the initial form request and first contact from the browser, and a request for more information on a message selected from the selection list (this message can ask for message data or descriptor). The application is able to tell what type of page it has by using hidden HTML fields and by checking the value of attributes in the CGI. The MQQueueB is then able to create the appropriate object for the request.

The browser object will then create an appropriate HTML page that is placed in a message and put onto the reply queue specified by the request.

This illustrates a multi-shot conversation between the client and server application. To enable this some sort of context needs to be supplied by the server application. The sample achieves this by the use of hidden fields in the HTML forms it sends back.

The sample consists of the following files:

- MQQueueB.html

This is the HTML form to start the session with the application.

- MQQueueB

 This is the executable for the application.

7.3.7.1 MQQueueB

The application is currently a long-running server that is started by:

MQQueueB *QueueName QueueManagerName*

The server can only be terminated by a kill.

7.3.7.2 MQQueueB.html

This HTML form has input fields to enter the name of the queue and queue manager that are to be browsed. It also has hidden fields that indicate what queue and queue manager the gateway will be using to put messages on for the application. The default setting of these is:

Queue MQQueueB.Queue

Queue manager Set to blank for the default queue manager

To use the sample, either create a queue of this name and use this and the default queue manager name to invoke MQQueueB, or edit the HTML to use any other queue and queue manager names required.

7.3.8 CGI Put Sample

This sample provides an HTML form with a queue and queue manager name entry fields along with a list box for message data. The button sends a POST request for the amqwput CGI program which then takes the CGI content and puts the message data onto the appropriate queue. This is essentially a CGI version of the MQSeries sample amqsput, shown on Figure 153 on page 315 and Figure 154 on page 321.

```
/*********************************************************************/
/*                                                                   */
/* MODULE NAME      amqwput0.cpp                                     */
/*                                                                   */
/* DESCRIPTIVE NAME Sample program that puts messages from           */
/*                  a message queue (example using MQPUT)            */
/*                  This is a modified version of the standard       */
/*                  MQSeries sample amqsput0.c  that allows          */
/*                  the pgm to be called by a CGI action and         */
/*                  write out the output in correct format           */
/*                                                                   */
/* Statement:       Licensed Materials - Property of IBM             */
/*                                                                   */
/*                  MA80 and MA81 SupportPac                         */
/*                  (c) Copyright IBM Corp. 1995.                    */
/*                                                                   */
/*                  See Copyright Instructions.                      */
/*                                                                   */
/*                  All rights reserved.                             */
/*                                                                   */
/*                  U.S. Government Users Restricted Rights - use,   */
/*                  duplication or disclosure restricted by GSA      */
/*                  ADP Schedule Contract with IBM Corp.             */
/*                                                                   */
/* Status:          Version 1 Release 1                              */
/* Genesis:         9th April 1996                                   */
/*                                                                   */
/*  NOTES :-                                                         */
/*    DEPENDENCIES = none                                            */
/*    RESTRICTIONS = none                                            */
/*    MODULE TYPE  = C++ source file                                 */
/*    PROCESSOR    = UNIX/PC                                         */
/*                                                                   */
 /*********************************************************************/
 /*                                                                   */
 /* Function:                                                         */
 /*                                                                   */
 /*                                                                   */
 /*   AMQWPUT0 is a sample C program to put messages on a message     */
 /*   queue, and is an example of the use of MQPUT.                   */
 /*                                                                   */
 /*      -- messages are sent to the queue named by the parameter     */
 /*                                                                   */
 /*      -- gets lines from StdIn, and adds each to target            */
 /*         queue, taking each line of text as the content            */
 /*         of a datagram message; the sample stops when a null       */
 /*         line (or EOF) is read                                     */
 /*                                                                   */
 /*      -- writes a message for each MQI reason other than           */
 /*         MQRC_NONE; stops if there is a MQI completion code        */
 /*         of MQCC_FAILED                                            */
 /*                                                                   */
```

Figure 153 (Part 1 of 6). C Program for the MQSeries Gateway

```
/*    Program logic:                                              */
/*         MQOPEN target queue for OUTPUT                         */
/*         while end of input file not reached,                   */
/*         .  read next line of text                              */
/*         .  MQPUT datagram message with text line as data       */
/*         MQCLOSE target queue                                   */
/*                                                                */
/*                                                                */
/******************************************************************/
/*                                                                */
/*   AMQWPUTO has 2 parameters                                    */
/*                 - the name of the target queue (required)      */
/*                 - queue manager name (optional)                */
/*                                                                */
/******************************************************************/
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
   /* includes for MQI */
#include <cmqc.h>
#include "CGIPart.h"
#include "CGIPartSet.h"
#include <stream.h>
#include "URLDecoder.h"
#ifdef __OS2__
#include <os2.h>
#endif

int main(int argc, char **argv)
{

  /*   Declare file and character for sample input                */
  FILE *fp;
  int    i;  /* auxiliary counter  */

  /*   Declare MQI structures needed                              */
  MQOD     od = {MQOD_DEFAULT};    /* Object Descriptor           */
  MQMD     md = {MQMD_DEFAULT};    /* Message Descriptor          */
  MQMD     mdDefault = {MQMD_DEFAULT};    /* Message Descriptor          */
  MQPMO   pmo = {MQPMO_DEFAULT};   /* put message options         */
     /** note, sample uses defaults where it can **/

  MQHCONN  Hcon;                   /* connection handle           */
  MQHOBJ   Hobj;                   /* object handle               */
  MQLONG   O_options;              /* MQOPEN options              */
  MQLONG   C_options;              /* MQCLOSE options             */
  MQLONG   CompCode;               /* completion code             */
  MQLONG   OpenCode;               /* MQOPEN completion code      */
  MQLONG   Reason;                 /* reason code                 */
  MQLONG   CReason;                /* reason code for MQCONN      */
  MQLONG   buflen;                 /* buffer length               */
  char     buffer[100];            /* message buffer              */
  char     QMName[50];             /* queue manager name          */
```

Figure 153 (Part 2 of 6). C Program for the MQSeries Gateway

```
unsigned int contentLength = 0;
long bLocalReturnCode = TRUE;
CGIPartSet the_CGIPartSet;
unsigned long ulNameLength;
 char *messageBuffer;
 char *line;
 URLDecoder theDecoder;

 if(strcmp(getenv("REQUEST_METHOD"),"POST")) {
     printf("This script should be referenced with a METHOD of POST.\n");
     printf("If you don't understand this, see this ");
     printf("<A HREF=\"http://www.ncsa.uiuc.edu/SDG/Software/Mosaic/Docs/\
            fill-out-forms/overview.html\">forms overview</A>.%c",10);
     exit(1);
     }
 if(strcmp(getenv("CONTENT_TYPE"),"application/x-www-form-urlencoded")) {
     printf("This script can only be used to decode form results. \n");
     exit(1);
     }

printf("Content-type: text/plain\n\n");

contentLength = atoi(getenv("CONTENT_LENGTH"));
messageBuffer = new char[contentLength +1];

cin.read(messageBuffer, contentLength);
messageBuffer[cin.gcount()] = '\0';

the_CGIPartSet.Initialize(messageBuffer, contentLength);

ulNameLength = MQ_Q_MGR_NAME_LENGTH;
bLocalReturnCode = the_CGIPartSet.getPartValue("PutQueueManager",
                                               QMName,
                                               &ulNameLength);

if (bLocalReturnCode == FALSE) {
   // No QM name use default ...
   QMName[0] = 0;    /* default */
   } /* endif */

ulNameLength = MQ_Q_NAME_LENGTH;
bLocalReturnCode = the_CGIPartSet.getPartValue("PutQueue",
                                               od.ObjectName,
                                               &ulNameLength);
if (bLocalReturnCode != TRUE) {
   // We had a duff form request come in
   printf("Required parameter missing - queue name\n");
   exit(99);
   } /* endif */

printf("Sample AMQSPUT0 start\n");
```

Figure 153 (Part 3 of 6). C Program for the MQSeries Gateway

```
/*****************************************************************/
/*                                                               */
/*   Connect to queue manager                                    */
/*                                                               */
/*****************************************************************/
MQCONN(QMName,                  /* queue manager                 */
       &Hcon,                   /* connection handle             */
       &CompCode,               /* completion code               */
       &CReason);               /* reason code                   */

/* report reason and stop if it failed     */
if (CompCode == MQCC_FAILED)
{
  printf("MQCONN ended with reason code %ld\n", CReason);
  exit(CReason);
}

/*****************************************************************/
/*                                                               */
/*   Use parameter as the name of the target queue               */
/*                                                               */
/*****************************************************************/
printf("target queue is %s\n", od.ObjectName);

/*****************************************************************/
/*                                                               */
/*   Open the target message queue for output                    */
/*                                                               */
/*****************************************************************/
O_options = MQOO_OUTPUT           /* open queue for output       */
       + MQOO_FAIL_IF_QUIESCING; /* but not if MQM stopping     */
MQOPEN(Hcon,                      /* connection handle           */
       &od,                       /* object descriptor for queue */
       O_options,                 /* open options                */
       &Hobj,                     /* object handle               */
       &OpenCode,                 /* MQOPEN completion code      */
       &Reason);                  /* reason code                 */

/* report reason, if any; stop if failed      */
if (Reason != MQRC_NONE)
{
  printf("MQOPEN ended with reason code %ld\n", Reason);
}

if (OpenCode == MQCC_FAILED)
{
  printf("unable to open queue for output\n");
}
```

Figure 153 (Part 4 of 6). C Program for the MQSeries Gateway

```
/*******************************************************************/
/*                                                                 */
/*   Read lines from the file and put them to the message queue    */
/*   Loop until null line or end of file, or there is a failure    */
/*                                                                 */
/*******************************************************************/
CompCode = OpenCode;        /* use MQOPEN result for initial test */
fp = stdin;

ulNameLength = sizeof(buffer);
bLocalReturnCode = the_CGIPartSet.getPartValue("PutData",
                                               buffer,
                                               &ulNameLength);
theDecoder.decodeInPlace(buffer);

line = strtok(buffer, "\n");

while (CompCode != MQCC_FAILED)
{
  if (line != NULL)
  {

   buflen = strlen(line);
  /*******************************************************************/
  /*                                                                 */
  /*   Put each buffer to the message queue                          */
  /*                                                                 */
  /*******************************************************************/
    md = mdDefault;
    memcpy(md.Format,          /* character string format          */
           MQFMT_STRING, MQ_FORMAT_LENGTH);
    MQPUT(Hcon,                /* connection handle                */
          Hobj,                /* object handle                    */
          &md,                 /* message descriptor               */
          &pmo,                /* default options (datagram)       */
          buflen,              /* buffer length                    */
          line,              /* message buffer                   */
          &CompCode,           /* completion code                  */
          &Reason);            /* reason code                      */

    /* report reason, if any */
    if (Reason != MQRC_NONE)
    {
      printf("MQPUT ended with reason code %ld\n", Reason);
    }
    line = strtok(NULL, "\r\n");
  }
  else   /* satisfy end condition when empty line is read */
    CompCode = MQCC_FAILED;
}
```

Figure 153 (Part 5 of 6). C Program for the MQSeries Gateway

```
  /*****************************************************************/
  /*                                                               */
  /*   Close the target queue (if it was opened)                   */
  /*                                                               */
  /*****************************************************************/
  if (OpenCode != MQCC_FAILED)
  {
    C_options = 0;                  /* no close options            */
    MQCLOSE(Hcon,                   /* connection handle           */
            &Hobj,                   /* object handle              */
            C_options,
            &CompCode,               /* completion code            */
            &Reason);                /* reason code                */

    /* report reason, if any     */
    if (Reason != MQRC_NONE)
    {
      printf("MQCLOSE ended with reason code %ld\n", Reason);
    }
  }

  /*****************************************************************/
  /*                                                               */
  /*   Disconnect from MQM if not already connected                */
  /*                                                               */
  /*****************************************************************/
  if (CReason != MQRC_ALREADY_CONNECTED)
  {
    MQDISC(&Hcon,                   /* connection handle           */
           &CompCode,               /* completion code             */
           &Reason);                /* reason code                 */

    /* report reason, if any     */
    if (Reason != MQRC_NONE)
    {
      printf("MQDISC ended with reason code %ld\n", Reason);
    }
  }

  /*****************************************************************/
  /*                                                               */
  /* END OF AMQWPUT0                                               */
  /*                                                               */
  /*****************************************************************/
  printf("Sample AMQWPUT0 end\n");
  delete [] messageBuffer;
  return(0);
}
```

Figure 153 (Part 6 of 6). C Program for the MQSeries Gateway

```
<HEAD>
<TITLE>MQSeries Internet Gateway Put Sample</TITLE>
</HEAD>

<BODY BGCOLOR="#E0E0FF">
<center>
<A HREF="./MQGate.html">
<img src="./images/MQPuts.gif" height=124 width=435
alt="• MQGate Page ‘" border=0>
</A>
<hr noshade size=1 width=545 align=center>
</center>
<FORM ACTION="/cgi-bin/amqwput" METHOD="POST">
<P>This is a sample frontend to do an MQPUT.
You will need a FORM capable browser.</P>
<p>Queue Manager: <INPUT NAME="PutQueueManager" VALUE=""></p>
<p>Queue: <INPUT NAME="PutQueue" VALUE=""></P>
<P>Enter the message data:</P>
<textarea NAME="PutData" cols="255" rows="20"></textarea>
</P>
<hr noshade size=1 width=545 align=center>
<P>
<font size=+1>
<INPUT TYPE="submit" VALUE="Put">
</font>
</P>
</FORM>
<hr noshade size=1 width=545 align=center>
</BODY>
</HTML>
```

Figure 154. HTML File for the MQSeries Gateway

7.3.9 CGI Get Sample

This sample provides an HTML form with a queue and queue manager name entry fields. The button sends a POST request for the amqwget CGI program which then takes the CGI content to open the appropriate queue, gets any messages on the queue and returns the data content back to the Web browser. This is essentially a CGI version of the MQSeries sample amqsget.

7.3.10 Application Programming Using the Gateway

This section shows you how to develop your ARP using the gateway.

7.3.10.1 Context Management

Currently all context management needs to be done by the application.

The application needs to either specify which queue and queue manager within any HTML is destined for the gateway or use the default values in the gateway ini file. See MQSeries Queue Manager and Queue. The application will probably also want to embed some of its own context information inside the HTML (for example, a page ID) so that when an application receives a message containing CGI data, it has some way of knowing where it came from and what to do with it.

7.3.10.2 Message Management

The application is also responsible for management of messages on the queues used by the gateway. There is the potential for unwanted messages to appear on the gateway reply queue (that is, a Web browser cancelled before a reply arrived). One way to deal with this is to set the expire time on any messages generated by the application.

7.3.11 Source Code

All the source code for the gateway and samples is available within the SupportPac. When expanded, a subdirectory source will be created, which has the following structure:

```
source/Makefile    - make file for whole of SupportPac
source/Makerule    - rules for all the make files
source/gateway     - directory containing server source code
source/samples     - directory containing all source for samples
source/cgilib      - directory containing source code for library
                     used by samples and gateway
source/bin         - target directory for all executables
source/log         - target directory for any logs created by make
```

The three subdirectories containing source code also have a `Makefile` in them. These individual make files are called by `Makefile` in the source directory. The default is to build an MQSeries Internet Gateway executable, MQGate, which uses the server version of MQSeries. The make file `source\gateway\Makefile` can be used to create a version using the client library. To do this, execute this make file with the command:

```
make MQGateClient
```

We use one command, rather than two separate executables, since the name of the CGI program (in this case MQGate) is referenced in all of the HTML forms in the samples.

The file `Makerule` may need editing to reflect the installation of MQSeries. There are variables in this file that need to be set to the correct paths.

7.3.11.1 Gateway Code

The gateway consists of the following classes, each class has a cpp and an h file:

MQGateway: This class encapsulates the MQSeries Internet Gateway. The post method will send the initial MQSeries message containing the CGI name/value pairs. Then it performs a get to obtain a response message, which it then sends back to the Web browser via stdout and the Web server. This class uses the formRequest and formResponse classes to create and access messages. It also uses the ConfigurationSet and ConfigurationPart classes.

formRequest: This class implements an object that transforms a CGI POST request string into an MQSeries message.

formResponse: This class implements an object that is used to transform an MQSeries message to stdout.

StatusPage: This class inherits from HTMLPage and provides an object that reads a file containing HTML and writes it to stdout, inserting some hidden fields that the gateway needs to process any action from the StatusPage.

In addition to the classes there are the following files:

MQGate.cpp: This is the main for the program MQGate that is invoked by the HTML page that contains a URL that points to the gateway.

7.3.11.2 Library Classes

The gateway and samples share a library of classes. These are in the source/cgilib:

CGIPart: This class is a name/value pair from the CGI POST string.

CGIPartSet: This class is a set of CGIParts. It is used to create a set of CGIPart objects, which is searchable, from a CGI Post request that has been read in from stdin.

ConfigurationPart: This class controls the configuration of the gateway using name/value pairs in an ini file.

ConfigurationSet: This class is a set of ConfigurationParts. It is used as the searchable interface into an ini file that contains a set of ConfigurationParts that are name/value pairs.

HTMLPage: This class is used to read an HTML page from a source file and output it to stdout for the Web server.

URLDecoder: This class provides a set of methods to aid in dealing with URL-encoded data strings.

7.3.11.3 Samples

These are all the samples available in the support pack.

amqwget0.cpp: Source code for amqwget executable.

amqwput0.cpp: Source code for amqwput executable.

MQHost.cpp: This is the source for the MQHost executable. It creates a QueueProcessor object and calls the GetContinually method.

queuepro.cpp and queuepro.hpp: This is the source for the QueueProcessor class. This class provides a simple interface to enable basic queue processing to be performed. After setting and starting, a Get or GetContinually can be performed and when a message is retrieved the method MessageProcessor is invoked. The MessageProcessor method in this instance retrieves the TCP hostname and puts this in a message on the reply queue of the message received. This class can be inherited from and this method should then be overridden to perform processing desired.

MQQueueB.cpp: This is the main for the sample MQQueueB executable, which provides a simple HTML queue browser. It creates a QueueScanner object and calls the GetContinually method.

queuescan.cpp and queuescan.hpp: This is the source for the QueueScanner class. This class inherits from QueueProcessor which provides basic queue processing functions. This class implements the messageProcessor method to create an HTML page dependent on the message retrieved. This is the core part of the queue browser

sample. The message retrieved will contain the name/value pairs from the hidden fields in the HTML form (there will be the queue and queue manager names and also the page name). The page name is used to tag the form that indirectly generated the message. The value of the page name and the setting of other form options is used to generate a new object, an HTMLBrowseProcessor, an HTMLMessageDescriptor or an HTMLMessageDataCharacterFormat.

hqmsgda.cpp and hqmsgda.hpp: This is the source for the HTMLMessageDataCharacterFormat class. This class inherits from QueueProcessor which provides services to get messages. This class builds an HTML page to contain the message data from a message found from a Get.

hqmsglist.cpp and hqmsglist.hpp: This is the source for the HTMLBrowseProcessor class. This class inherits from QueueProcessor which provides services to get messages. This class builds an HTML page that contains a list of messages on the queue.

hqmsgmd.cpp and hqmsgmd.hpp: This is the source for the HTMLMessageDescriptor class. This class inherits from QueueProcessor which provides services to get messages. This class builds an HTML page to contain the message descriptor fields from a message found from a Get.

cache.cpp cache.hpp: This is the source for the Cache class, used by the MQHost and MQQueueB samples. This class provides a simple memory cache object.

7.4 AS400 Web Server Screen Translator

Most Web servers today require that you write scripts or programs to create interactive forms and applications for the World Wide Web. For most software providers, this can mean learning new tools and procedures if they want to support the World Wide Web. This is not true for AS/400 customers. With the AS/400 HTML Gateway function in WebConnection for OS/400, your current development tools work for creating WWW applications. Once your WWW applications are created, you can start using the Internet′s worldwide reach to open new marketing opportunities. Even existing AS/400 applications can run over the Web without modifying any code. There is no conversion program to run. Just install and configure WebConnection for OS/400, and the applications on your AS/400 system are ready to go.

So how does IBM do it?

AS/400 applications are inherently display-oriented. This means that each application creates a series of displays for use in its application. These displays are normally sent out in a 5250 data stream to the workstation or emulator, which shows the text. WebConnection for OS/400 intercepts this 5250 data stream and converts it to HTML, a language the Web understands. Any Web browser used for accessing the World Wide Web can work with the application.

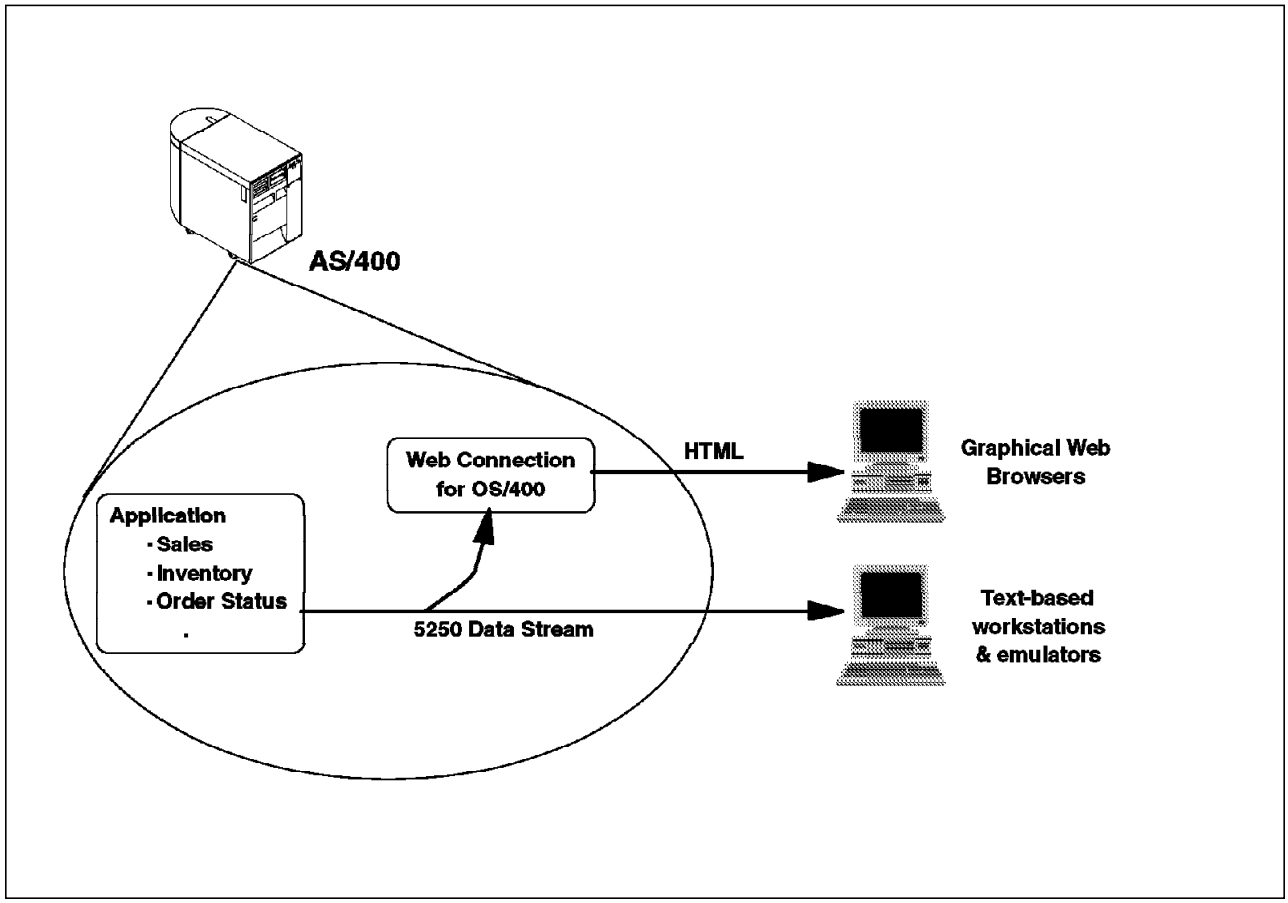


Figure 155. 5250 HTML Gateway

WebConnection for OS/400 means your business does not need to rely on one specific client platform. Any PC that has a Web browser installed can run AS/400 applications. There is no additional connection configuration. Just point your Web browser to the AS/400 system, and you are in business.

If your business writes AS/400 applications, then WebConnection for OS/400 means a wealth of new applications on the Internet. You do not need to retrain your programmers. They can continue using their existing development tools (RPG, COBOL, and DDS). Also, with AS/400 HTML Gateway in WebConnection for OS/400, your programmers can jazz up your applications by adding graphics. It requires only a small change to the DDS specifications, and it does not affect your workstation users.

Now that we know what a 5250 HTML gateway does, let's see some examples of the translation from text-based 5250 panels to something a Web client can see and use. For this, we are going to show you some OS/400 displays that have been translated to HTML by an early version of the workstation gateway support. The final look and feel may be quite different from what we will show you here.

1. Sign-on

 Figure 156 on page 327 shows a portion of the traditional AS/400 sign-on display converted now to HTML and displayed on a WebExplorer client. Note the functionality is really no different than with a normal text-based 5250 emulator.

The URL that your Web client needs to specify to evoke the 5250 to HTML WorkStation Gateway support will look something like this:

```
http://hostname:5061/WSG
```

Where

http: The WorkStation Gateway uses the HTTP protocol.

hostname This identifies the system to which the request will go. This could be just the host name or the fully qualified host name with domain.

:5061 5061 is the default well-known port for the WorkStation Gateway server. You must specify this port as your Web client will try to connect to port 80 by default if you fail to override this.

?exit_information

Not shown in the above example are the optional parameters that can be used to pass information from the client to the WorkStation Gateway server running on the AS/400. Characters following the WSG will be interpreted as parameters to be passed to the server job. For the initial connection, these parameters could be a user ID and password used to direct the new client directly to a 5250 application without the need to sign on to the AS/400. Later, after the session has been established, what follows after the WSG is information to allow the AS/400 to route this screen to the proper WorkStation Gateway server. This is because the AS/400 must save state, while using a protocol like HTTP which does *not* save state. Look closely at the bottom of all the figures in this section for the URL used to save state.

Please see 7.4.2, “5250 HTML Workstation Gateway Application Logon Exit Program” on page 330 for more information about the WorkStation Gateway exit program.

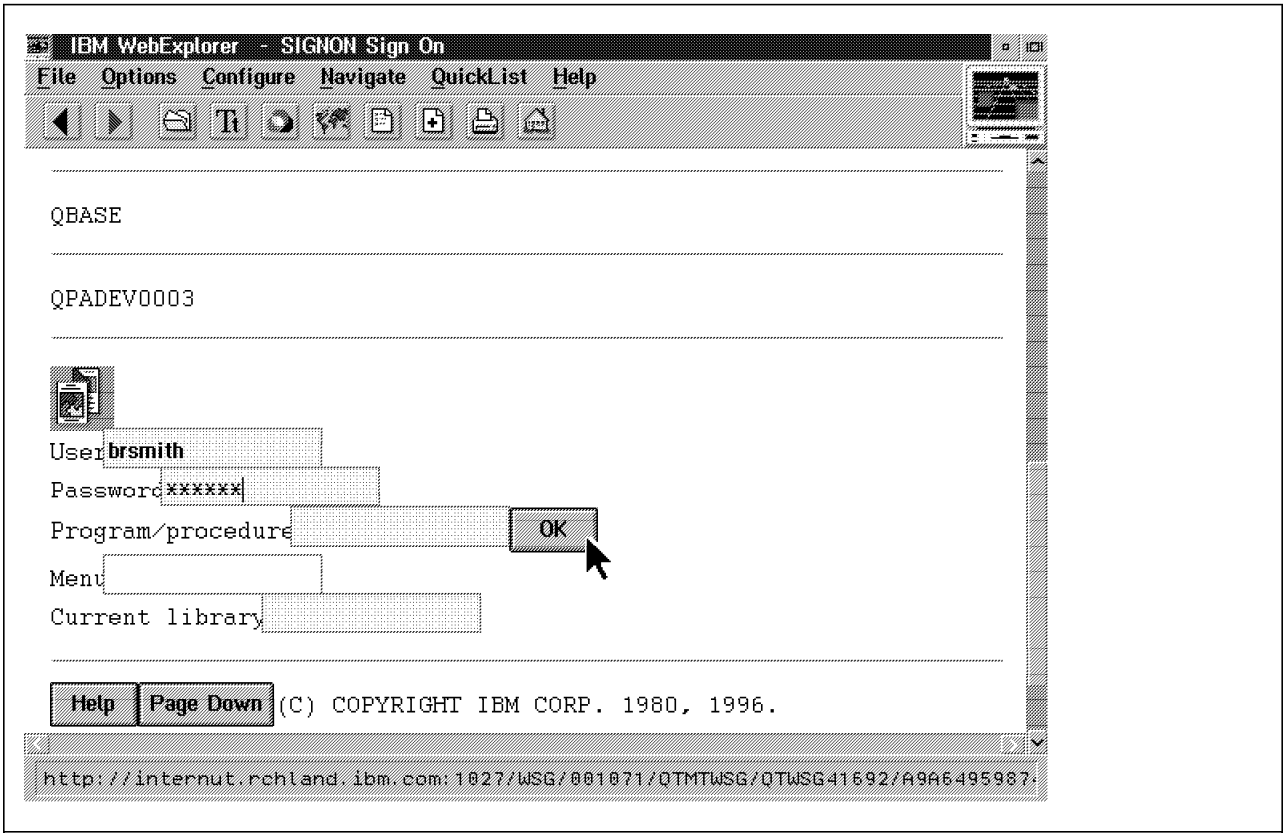


Figure 156. A Portion of the AS/400 Sign-On as Seen by the Workstation Gateway

2. Command Entry

 Figure 157 on page 328 shows the Command Entry display for the WebExplorer client. For example, the Functions list allows you to retrieve the previous command.

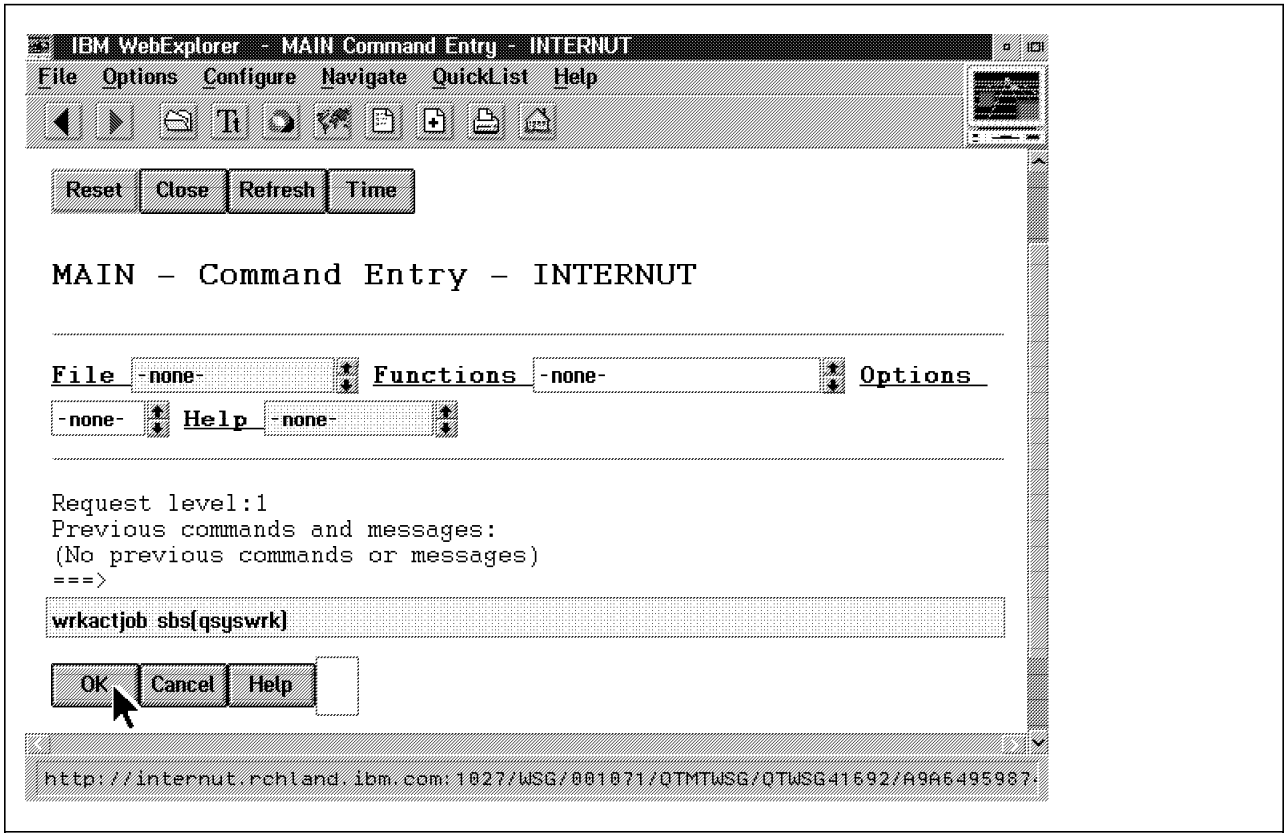


Figure 157. The Command Entry Display as Seen by the Workstation Gateway

3. Work Active Job

 Figure 158 on page 329 shows the Work with Active Job display for the WebExplorer client. Note that your Web client must be able to display tables. You select the job with the check box, and then select the function you want to perform on that job. As shown, two jobs have been selected and the mouse pointer is poised to select the **Work with** function.

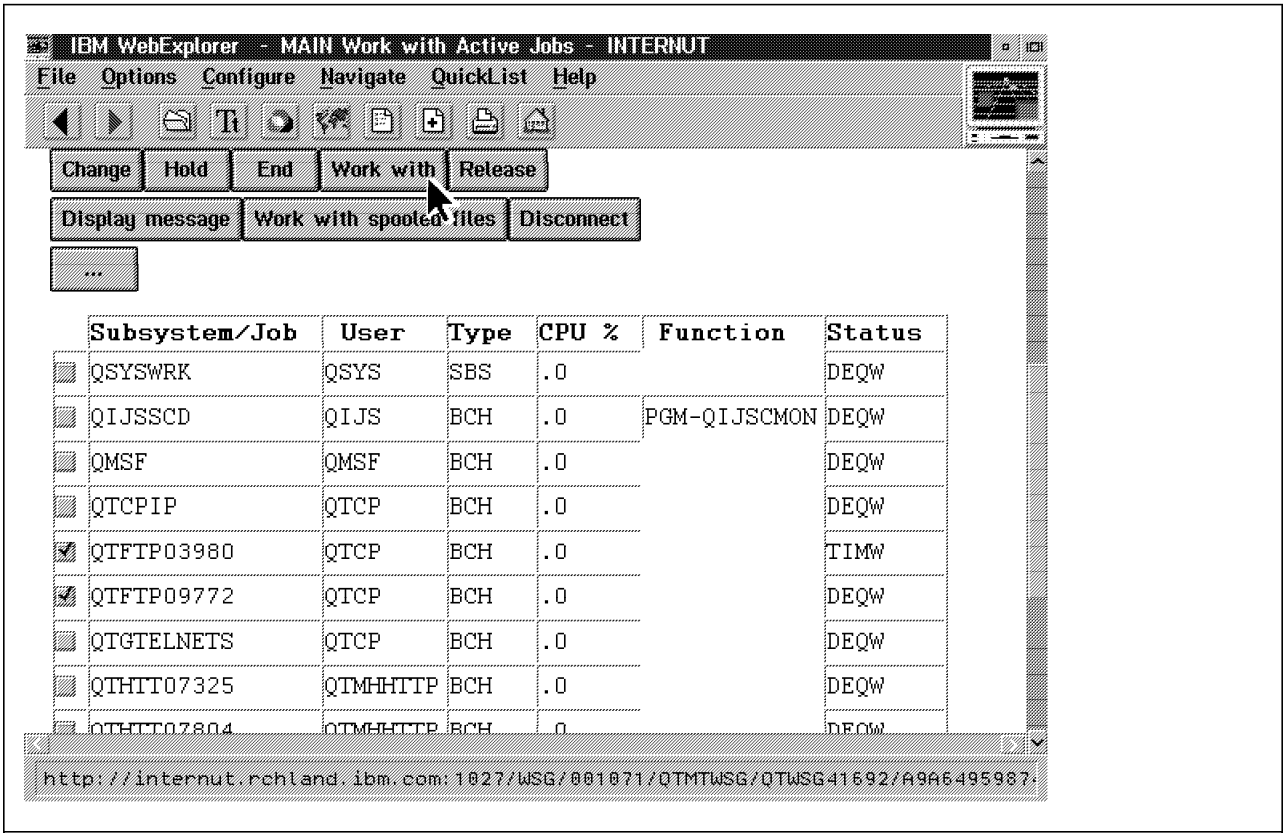


Figure 158. The Work with Active Job Display As Seen by the Workstation Gateway

7.4.1 The 5250 HTML Gateway Server

Do not confuse the 5250 HTML Gateway with the HTTP Web Server. The HTTP Web Server allows the AS/400 system to act as a WWW server in the Internet. The 5250 HTML Gateway converts your 5250 data stream to HTML. Both can be started and function independent from each other.

The 5250 HTML Gateway is a TCP/IP application that services requests from HTTP clients. After the initial request is received from a client, that client is considered ″active″ and all future connections requests for that client occur over an arbitrary port number.

The client remains active until the session is signed off or an inactivity timeout limit is reached.

Note

The 5250 HTML Gateway maintains the illusion that the browser is logically connected to the AS/400 system even though every transaction between the browser and the AS/400 server is disconnected. The AS/400 server maintains the virtual terminal API connection indefinitely or until the browser logs off or the inactivity timeout value is exceeded.

The 5250 HTML server is started through the following command:

```
STRTCPSVR SERVER(*WSG)
```

and ended with the following:

```
ENDTCPSRV SERVER(*WSG)
```

Alternatively, it is started through the AUTOSTART option of the STRTCP command. The jobs are named QTVQVTnnnnn where nnnnn is a unique numeric string that is derived from the time stamp.

The format of a link in an HTML document is called a Universal Resource Locator (URL). For HTTP, the URL identifies the protocol that the browser should use when contacting the server (for example, HTTP, FTP, WAIS, Gopher, and so on) and the location of the server, and of the requested object. HTTP has the following form:

```
http://hostname:port/path
```

The port numbers for most TCP/IP applications such as FTP, Telnet, or WWW are predefined or you might say *well-known* numbers, which means everyone knows them and uses the same port numbers.

The 5250 HTML Gateway does not have such a well-known port number such as the HTTP server has. Therefore, the port number used by the AS/400 Virtual Terminal Gateway is found by querying the local TCP/IP configuration services database. To establish a 5250 HTML Gateway session, you must connect using the form:

```
http://hostname:port
```

where port is the configured port number for that 5250 HTML Gateway. The default is a TCP port of 5061.

The 5250 server is organized into the following:

- A single *parent* job that *listens* and *accepts* connections from HTTP browser clients. It is important to note that the port used by 5250 HTML Gateway is different from the port of the HTTP Server because the 5250 HTML Server is a new type of server for which there is no well-known port. The parent job has only one function to hand off connection requests to child jobs.
- One or more *child* jobs. A child job performs the actual work to satisfy the client connect request.

This technique allows you to do a multiplexing of connections within a single batch job.

7.4.2 5250 HTML Workstation Gateway Application Logon Exit Program

An application logon exit program (QAPP0100) will allow bypassing the AS/400 sign-on display and invoking an application program directly without the client browser having to send a user profile or password. This allows the customer the option of providing *any* application to client browsers without requiring a sign on. This is done by calling a customer program that authenticates the client request and provides sign-on information to the 5250 HTML Gateway Server.

The 5250 HTML Gateway Server uses the output of the customer's User Exit as input to the Virtual Terminal APIs and performs the sign-on action on behalf of the client browser.

When the user exit is given control, it must perform any desired validation using the supplied Internet Protocol address and any of the supplied operation specific information extracted after the /WSG string in the URL. Setting the Allow Operation output determines whether the automatic logon is performed, or whether an error message is returned to the client browser.

If the operation is allowed, then the user exit must return the user profile, password, current library, and program. All output must be non-NULL or else an error is returned to the client browser.

7.4.3 Configure TCP/IP Workstation Gateway (CFGTCPWSG) Main Menu

The easiest way to configure the 5250 HTML Gateway is to use the menus. The following examples show the sequence of the configuration commands.

The following display appears if CFGTCPWSG is entered from the command line, or if CFGTCPAPP option 15 is selected.

```
                  Configure TCP/IP Workstation Gateway
                                                          System:  SYSNMO11
  Select one of the following:

      1. Change workstation gateway attributes

   Related options:
     10. Configure HTTP
     11. Work with autoconfigure virtual devices
     12. Work with limit security officer device access

 Selection or command
 ===> ____________________________________________________________________

 F3=Exit   F4=Prompt   F9=Retrieve   F12=Cancel
```

Figure 159. CFGTCPWSG Display

- Option 1 - Prompts the CHGWSGA CL command.
- Option 10 - Calls the CFGTCPHTTP CL command.
- Option 11 - Calls WRKSYSVAL SYSVAL(QAUTOVRT)
- Option 12 - Calls WRKSYSVAL SYSVAL(QLMTSECOFR)

7.4.4 Change Workstation Gateway Attributes (CHGWSGA) CL Command Prompt

The following display appears if the CHGWSGA CL command is prompted from the command line or if CFGTCPWSG option 1 is selected.

The values shown are the current values as determined by the Prompt Override Program for CHGWSGA.

```
      Change Workstation Gateway Attributes (CHGWSGA)
                                                        System:  SYSNM011
Type choices, press Enter.

Autostart  . . . . . . . . . . .   *NO           *NO, *YES, *SAME
Number of clients per server job   20           1-50, *SAME, *DFT
Inactivity timeout . . . . . . .   10           0-60 minutes, *SAME, *DFT
Data request timeout . . . . . .   10           1-1200 seconds, *SAME, *DFT
Special key placement  . . . . .   *TOP          *TOP, *BOT, *SAME
Function key placement . . . . .   *BOT          *BOT, *TOP, *SAME
Top banner URL . . . . . . . . .   *NONE_____________________________________
_____________________________________________________________________________
Bottom banner URL  . . . . . . .   *NONE_____________________________________
_____________________________________________________________________________
_____________________________________________________________________________
_____________________________________________
Coded character set identifier     00819        1-65533, *SAME, *DFT

F3=Exit   F4=Prompt   F5=Refresh   F12=Cancel   F13=How to use this display
F24=More keys
```

Figure 160. Change Workstation Gateway Attributes Display

7.4.4.1 Timeout Values

Since many clients can be expected to use the 5250 HTML Gateway Server, it is important to always try to have free servers waiting for new connect requests. To stay ahead of potential load demands, jobs are *pre-started* to avoid SBMJOB latency when a new server job is close to being needed.

When we say pre-started, we mean that we submit a new child server with the SBMJOB when the number of available jobs goes below threshold limits (remember we are multiplexing connections within a single batch server job). The threshold limit is determined based upon the value selected for the configured number of clients.

We have two types of timeouts for the 5250 HTML Gateway Server:

1. Inactivity timeout (INACTTIMO) - default 10 minutes

 Specifies the number of minutes the system allows a Workstation Gateway session to remain inactive before it is ended. When a WSG session is inactive longer than the specified length of time, it is ended.

 Note: It may take the system an additional 1 to 120 seconds to end the inactive session.

2. Data request timeout (DTARQSTIMO) - default 10 seconds

 Specifies the number of seconds the system allows a Workstation Gateway session to wait from the time a Workstation Gateway client requests data to the time the data is sent by the Workstation Gateway server job.

Both timeout values can be changed in the CHGWSGA command.

7.4.4.2 What Happens with My Existing Display Files?

Your existing display files need not be changed. You can use all DDS specifications as you did before. The DDS becomes (when compiled) a 5250 data stream. This means that the DDS keywords such as DSPATR(UL), BLINK, CHECK, and so on are translated in a coded string of data. In this data string, each field is preceded by one or more attribute bytes. This information makes a field such as a customer name underlined, protected, or blinking.

The AS/400 system (or more precise, the twinax workstation IOP (input/output processor)) sends out this generated 5250 data stream to your ″green″ 5250 screen. The hardware of your screen then interprets this stream of data and produces a protected, underlined, or blinking field on your display.

This is the way it works today. With V3Rx, the 5250 HTML gateway intercepts this 5250 data stream and converts it ″on the fly″ to an HTML data stream. Let′s look at an example to make it more comprehensive.

First, we show you a simple DDS example of a display and how it looks on a 5250 workstation (green screen).

Note: This DDS example is not using any new techniques or HTML keywords.

```
A                                      DSPSIZ(24 80 *DS3)
A         R RECORD1
A                                   3 18'Display of the customer master rec-
A                                       ord'
A                                       DSPATR(UL)
A                                   3 62' '
A                                   6  9'CUSNUM:'
A           CUSNUM    R        B  6 18REFFLD(CUSREC/CUSNUM QIWS/QCUSTCDT)
A                                   8  9'LSTNAM:'
A           LSTNAM    R        B  8 18REFFLD(CUSREC/LSTNAM QIWS/QCUSTCDT)
A                                  10  9'STREET:'
A           STREET    R        B 10 18REFFLD(CUSREC/STREET QIWS/QCUSTCDT)
A                                  12  9'ZIPCOD:'
A           ZIPCOD    R        B 12 18REFFLD(CUSREC/ZIPCOD QIWS/QCUSTCDT)
A                                  14 11'CITY:'
A           CITY      R        B 14 18REFFLD(CUSREC/CITY QIWS/QCUSTCDT)
A                                  18  9'BALDUE:'
A           BALDUE    R        B 18 18REFFLD(CUSREC/BALDUE QIWS/QCUSTCDT)
```

Figure 161. DDS Source for Our Customer Master Record

The preceding DDS looks the same as this on a 5250 display station:

```
                Display of the customer master record

        CUSNUM:

        LSTNAM:

        STREET:

        ZIPCOD:

          CITY:


        BALDUE:
```

Figure 162. Customer Master Record DDS on the Traditional Text 5250 Display

Now let's see what the 5250 HTML Gateway made out of our DDS specifications. The following display shows the result of the 5250 data stream conversion process. Note that this does not mean that you had to recompile the display file. The 5250 HTML Gateway did this automatically "on the fly" for you. When the 5250 HTML Gateway detected that the terminal that receives the 5250 data stream was a virtual terminal (that is, a PC), the 5250 data stream was converted to the HTML data stream.

```
<BR>Display of the customer master record
<BR>CUSNUM:<INPUT TYPE="TEXT" NAME="afield.006-018" VALUE="" SIZE=7 MAXLENGTH=7>
<BR>LSTNAM:<INPUT TYPE="TEXT" NAME="afield.008-018" VALUE="" SIZE=8 MAXLENGTH=8>
<BR>STREET:<INPUT TYPE="TEXT" NAME="afield.010-018" VALUE="" SIZE=13 MAXLENGTH=13>
<BR>ZIPCOD:<INPUT TYPE="TEXT" NAME="afield.012-018" VALUE="" SIZE=6 MAXLENGTH=6>
<BR>CITY:<INPUT TYPE="TEXT" NAME="afield.014-018" VALUE="" SIZE=6 MAXLENGTH=6>
<BR>BALDUE:<INPUT TYPE="TEXT" NAME="afield.018-018" VALUE="" SIZE=7 MAXLENGTH=7>
```

Figure 163. HTML Automatically Generated by the 5250-HTML Gateway

Finally, let's see how this looks on an OS/2 Web browser.

Note: The result you see on a Web browser is totally dependent upon how you configured the browser. If you choose another font, another background color, or another font size, the actual appearance of your HTML data stream on your PC might look quite different from our example.

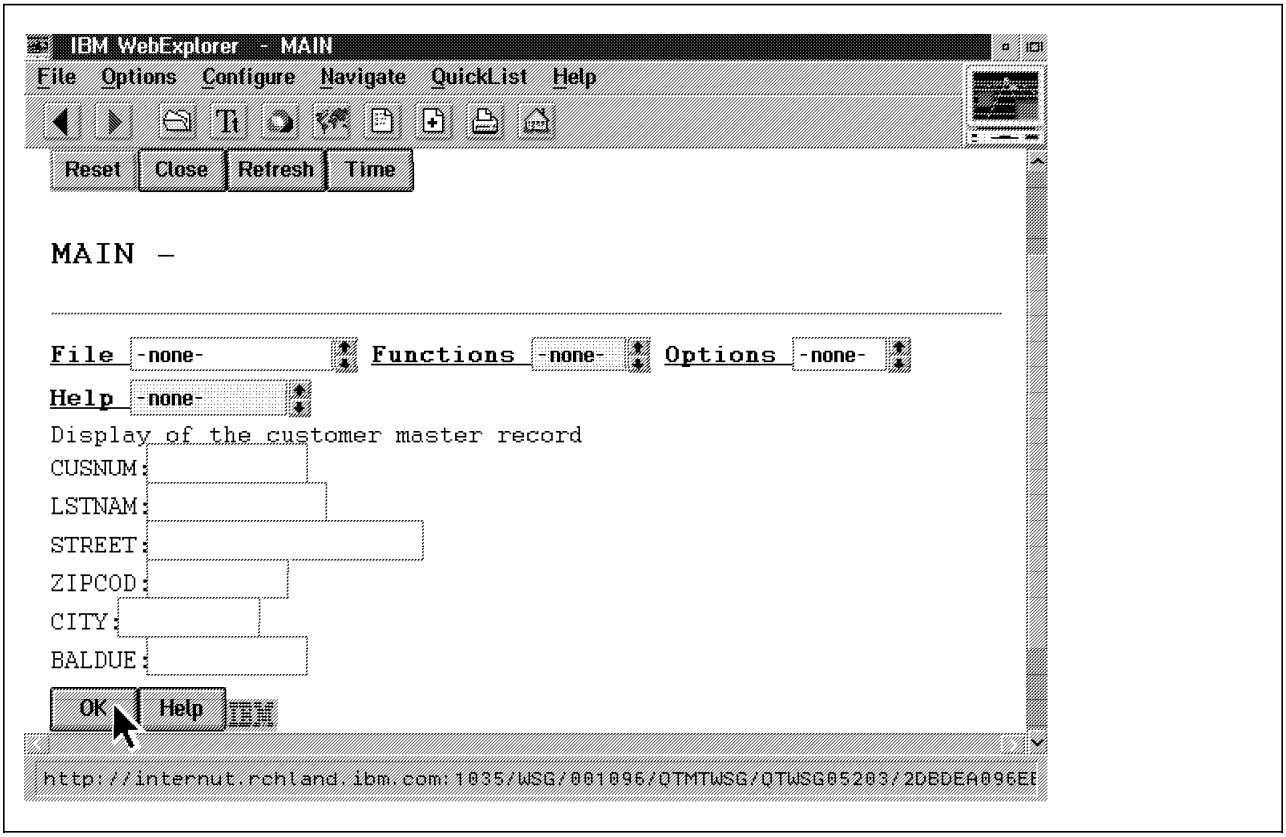


Figure 164. Display of Customer Master Record through 5250-HTML Gateway

7.4.5 How Can I Use the HTML Support for New Possibilities?

The 5250 HTML Gateway support allows the insertion of HTML tags into the DDS of a display file. This allows us to utilize the graphic capabilities of a Web browser with only minor changes to the existing DDS. For example, a customer can add graphics through the IMG HTML tag to an existing display file and display a graphic image along with the display.

Note: These HTML tags are only inserted into the data stream that flows to a terminal if the device query indicates that the device is a PC (or more precisely, an AS/400 5250 Workstation Gateway virtual terminal). Otherwise, the HTML tags are ignored for normal displays.

This simplifies and eases the handling of display files because only *one* source is needed for graphical workstations (that is, PCs) and green screens.

7.4.5.1 The New DDS Keyword

There is a new DDS keyword: HTML (HyperText Markup Language). This field level keyword can be treated the same as a usual constant. Two things are different from a common constant. First, you have to put the new keyword HTML before the constant, and second, the ″constant″ itself must consist of an HTML string that must use the HTML syntax.

Let′s take a look at a DDS example with HTML statements.

```
+... 1 ...+... 2 ...+... 3 ...+... 4 ...+... 5 ...+... 6 ...+... 7 ...+..
 A*----------------------------------------------------------------------
 A          R RCD1
 A  42                                   PUTOVR
 A  43                                   OVRDTA
 A                                    1  5'Regular DDS text' DSPATR(RI)
 A                                    3  3HTML('<html>')
 A                                    3  3HTML('<head>')
 A                                    3  3HTML('<title>Test Screen</title>'
```

Figure 165. Sample 5250 DDS Enhanced with the HTML Tag

Note: The plain text is mixed with so-called HTML tags.

What are HTML Tags?

HTML documents consists of plain text interspersed with markup commands called *tags*. The tags are instructions to the browser software on how to display the text. They are represented by strings enclosed in <angle brackets> the same as the words before.

Another thing to mention is that in the preceding example, ″normal″ DDS keywords and HTML specs are used within one source.

HTML is a tag language where the order of the tags determines when they are processed. Row and column have no meaning in such a tag language. In this case, the row and column are used to determine the order in which the HTML tags are sent to the browser.

With the HTML keyword, constant fields that have the same row and column value are processed in the order in which they appear in the DDS source.

How to Determine if HTML is Processed?

On the CRTDSPF command, the ENHDSP (enhanced display) parameter is used to ignore or process the HTML keywords. This setting can be changed dynamically.

7.4.5.2 Format of the HTML Specification

The new HTML specification can have two formats:

- HTML (datastring with a valid HTML tag)
- HTML (program-to-system-field)

A parameter is required after an HTML keyword. This parameter can be a valid HTML tag enclosed in single quotes, or a program variable. The program-to-system field can be any legal length and has to be alphanumeric (A in position 35).

Note: The syntax of the HTML tag is *not* syntax checked by the DDS compiler. The browser that receives the HTML sequence performs syntax checking.

7.4.5.3 Limitations/Restrictions

The following keywords are not allowed with the HTML keyword:

- COLOR
- DATE
- DFT
- DSPATR
- EDTCDE
- EDTWRD
- HLPID
- MSGCON
- NOCCSID
- OVRATR
- PUTRETAIN
- SYSNAME
- TIME
- USER

The HTML keyword is not allowed on a field in a subfile record.

Chapter 8. Security on the Internet

The world of computers has changed dramatically over the past twenty-five years. Twenty-five years ago, most computers were centralized and managed by data centers. Computers were kept in locked rooms and staffs of people made sure they were carefully managed and physically secured. Links outside a site were unusual. Computer security threats were rare, and were basically concerned with insiders: authorized users misusing accounts, theft and vandalism, and so forth. These threats were well understood and dealt with using standard techniques: computers behind locked doors, and accounting for all resources. Computing in the 1990s is radically different. Many systems are in private offices and labs, often managed by individuals or persons employed outside a computer center and the big problem is systems connected into the Internet. With worldwide Internet connections, someone could get into your system from the other side of the world and steal your password in the middle of the night when your building is locked up. Viruses and worms can be passed from machine to machine.

The Internet allows the electronic equivalent of the thief who looks for open windows and doors; now a person can check hundreds of machines for vulnerabilities in a few hours. System administrators and decision makers have to understand the security threats that exist, what the risk and cost of a problem would be, and what kind of action they want to take to prevent and respond to security threats. Setting security policies and procedures really means developing a plan for how to deal with computer security. You need to first:

- Look at what you are trying to protect.
- Look at what you need to protect it from.
- Determine how likely the threats are.
- Implement measures which will protect your assets in a cost-effective manner.
- Review the process continuously, and improve things every time a weakness is found.

One old truism in security is that the cost of protecting yourself against a threat should be less than the cost recovering if the threat were to strike you. We can divide Internet security in a two diferent parts:

- Policies
- Technologies

The policies are theorical procedures. If these procedures are corretly used, the security can be improved and the possibilities of a security fail, reduced. The technologies are resources that use hardware and software to provide high levels of security, like firewalls and cryptograpy technics. But you can get the most efficient model of security only using the policies and the technologies together.

© Copyright IBM Corp. 1996

8.1 Policies

8.1.1 Organization Issues

The goal in developing an official site policy on computer security is to define the organization′s expectations of proper computer and network use and to define procedures to prevent and respond to security incidents. In order to do this, aspects of the particular organization must be considered. First, the goals and direction of the organization should be considered. For example, a military base may have very different security concerns from a those of a university. Second, the site security policy developed must conform to existing policies, rules, regulations and laws that the organization is subject to. Therefore it will be necessary to identify these and take them into consideration while developing the policy. Third, unless the local network is completely isolated and standalone, it is necessary to consider security implications in a more global context. The policy should address the issues when local security problems develop as a result of a remote site as well as when problems occur on remote systems as a result of a local host or user.

8.1.2 Who Makes the Policy?

Policy creation must be a joint effort by technical personnel, who understand the full ramifications of the proposed policy and the implementation of the policy, and by decision makers who have the power to enforce the policy. A policy that is neither implementable nor enforceable is useless. Since a computer security policy can affect everyone in an organization, it is worth taking some care to make sure you have the right level of authority in on the policy decisions. Though a particular group (such as a campus information services group) may have responsibility for enforcing a policy, an even higher group may have to support and approve the policy.

8.1.3 Who Is Involved?

Establishing a site policy has the potential for involving every computer user at the site in a variety of ways. Computer users may be responsible for personal password administration. Systems managers are obligated to fix security holes and to oversee the system. It is critical to get the right set of people involved at the start of the process. There may already be groups concerned with security who would consider a computer security policy to be their area. Some of the types of groups that might be involved include auditing/control, organizations that deal with physical security, campus information systems groups, and so forth. Asking these types of groups to ″buy in″ from the start can help facilitate the acceptance of the policy.

8.1.4 Responsibilities

A key element of a computer security policy is making sure everyone knows their own responsibility for maintaining security. A computer security policy cannot anticipate all possibilities; however, it can ensure that each kind of problem does have someone assigned to deal with it. There may be levels of responsibility associated with a policy on computer security. At one level, each user of a computing resource may have a responsibility to protect his or her account. Users who allow their account to be compromised increase the chances of compromising other accounts or resources. System

managers may form another responsibility level: they must help to ensure the security of the computer system. Network managers may reside at yet another level.

8.1.5 Risk Assessment

One of the most important reasons for creating a computer security policy is to ensure that efforts spent on security yield cost-effective benefits. Although this may seem obvious, it is possible to be mislead about where the effort is needed. As an example, there is a great deal of publicity about intruders on computers systems; yet most surveys of computer security show that for most organizations, the actual loss from ″insiders″ is much greater.

Risk analysis involves determining what you need to protect, what you need to protect it from, and how to protect it. It is the process of examining all of your risks, and ranking those risks by level of severity. This process involves making cost-effective decisions on what you want to protect. The old security adage says that you should not spend more to protect something than it is actually worth.

8.1.5.1 Identifying the Assets

One step in a risk analysis is to identify all the things that need to be protected. Some things are obvious, like all the various pieces of hardware, but some are overlooked, such as the people who actually use the systems. The essential point is to list all things that could be affected by a security problem, like:

- **Hardware:** Cpus, boards, keyboards, terminals, workstations, personal computers, printers, disk drives, communication lines, terminal servers, routers.
- **Software:** Source programs, object programs, utilities, diagnostic programs, operating systems, communication programs.
- **Data:** During execution, stored online, archived offline, backups, audit logs, databases, in transit over communication media.
- **People:** Users, people needed to run systems.
- **Documentation:** On programs, hardware, systems, local administrative procedures.
- **Supplies:** Paper, forms, ribbons, magnetic media.

8.1.5.2 Identifying the Threads

Once the assets requiring protection are identified, it is necessary to identify the threats to those assets. The threats can then be examined to determine what potential for loss exists. It helps to consider the threats you are trying to protect your assets from.

The following sections describe a few of the possible threats.

Unauthorized Access: A common threat that concerns many sites is unauthorized access to computing facilities. Unauthorized access takes many forms. One means of unauthorized access is the use of another user′s account to gain access to a system. The use of any computer resource without prior permission may be considered unauthorized access to computing facilities. The seriousness of an unauthorized access will vary

from site to site. For some sites, the mere act of granting access to an unauthorized user may cause irreparable harm by negative media coverage. For other sites, an unauthorized access opens the door to other security threats. In addition, some sites may be more frequent targets than others; hence the risk from unauthorized access will vary from site to site. The Computer Emergency Response Team (CERT), has observed that well-known universities, government sites and military sites seem to attract more intruders.

Disclosure of information: Another common threat is disclosure of information. Determine the value or sensitivity of the information stored on your computers. Disclosure of a password file might allow for future unauthorized accesses. A glimpse of a proposal may give a competitor an unfair advantage. A technical paper may contain years of valuable research.

Denial of service: Computers and networks provide valuable services to their users. Many people rely on these services in order to perform their jobs efficiently. When these services are not available when called upon, a loss in productivity results. Denial of service comes in many forms and might affect users in a number of ways. A network may be rendered unusable by a rogue packet, jamming, or by a disabled network component. A virus might slow down or cripple a computer system. Each site should determine which services are essential, and for each of these services determine the affect to the site if that service were to become disabled.

8.1.6 Policy Issues

There are a number of issues that must be addressed when developing a security policy. These are:

- Who is allowed to use the resources?
- What is the proper use of the resources?
- Who may have system administration privileges?
- What are the user's rights and responsibilities?
- What do you do with sensitive information?
- What happens when the policy is violated?

These issues are discussed below. In addition you may wish to include a section in your policy concerning ethical use of computing resources.

8.1.6.1 Who Is Allowed to Use the Resources?

One step you must take in developing your security policy is defining who is allowed to use your system and services. The policy should explicitly state who is authorized to use what resources.

8.1.6.2 What Is the Proper Use of the Resources?

After determining who is allowed access to system resources it is necessary to provide guidelines for the acceptable use of the resources. You may have different guidelines for different types of users (that is, students, faculty, external users). The policy should state what is acceptable use as well as unacceptable use. It should also include types of use that may be restricted. Define limits to access and authority. You will need to consider the level of access various users will have and what resources will be available or restricted to various groups of people. Your acceptable use policy should

clearly state that individual users are responsible for their actions. Their responsibility exists regardless of the security mechanisms that are in place. It should be clearly stated that breaking into accounts or bypassing security is not permitted.

The following points should be covered when developing an acceptable use policy:

- Is breaking into accounts permitted?
- Is cracking passwords permitted?
- Is disrupting service permitted?
- Should users assume that a file being world-readable grants them the authorization to read it?
- Should users be permitted to modify files that are not their own even if they happen to have write permission?
- Should users share accounts?

The answer to most of these questions will be no.

You may wish to incorporate a statement in your policies concerning copyrighted and licensed software. Licensing agreements with vendors may require some sort of effort on your part to ensure that the license is not violated. In addition, you may wish to inform users that the copying of copyrighted software may be a violation of the copyright laws and is not permitted.

Specifically concerning copyrighted and/or licensed software, you may wish to include the following information:

- Copyrighted and licensed software may not be duplicated unless it is explicitly stated that you may do so.
- Methods of conveying information on the copyright/licensed status of software.
- When in doubt, don't copy.

Your acceptable use policy is very important. A policy that does not clearly state what is not permitted may leave you unable to prove that a user violated the policy.

There are exception cases such as tiger teams and users or administrators wishing for licenses to hack, you may face the situation where users will want to hack on your services for security research purposes. You should develop a policy that will determine whether you will permit this type of research on your services and if so, what your guidelines for such research will be.

Points you may wish to cover in this area:

- Whether it is permitted at all.
- What type of activity is permitted: breaking in, releasing worms, releasing viruses, etc.
- What type of controls must be in place to ensure that it does not get out of control (separate a segment of your network for these tests).

- How you will protect other users from being victims of these activities, including external users and networks.
- The process for obtaining permission to conduct these tests.

In cases where you do permit these activities, you should isolate the portions of the network that are being tested from your main network. Worms and viruses should never be released on a live network.

You may also wish to employ, contract, or otherwise solicit one or more people or organizations to evaluate the security of your services, of which may include hacking. You may wish to provide for this in your policy.

8.1.6.3 Who May Have System Administration Privileges?

One security decision that needs to be made very carefully is who will have access to system administrator privileges and passwords for your services. Obviously, the system administrators will need access, but inevitably other users will request special privileges. The policy should address this issue. Restricting privileges is one way to deal with threats from local users. The challenge is to balance restricting access to these to protect security while giving people who need these privileges access so that they can perform their tasks. One approach that can be taken is to grant only enough privilege to accomplish the necessary tasks.

Additionally, people holding special privileges should be accountable to some authority and this should also be identified within the site′s security policy. If the people you grant privileges to are not accountable, you run the risk of losing control of your system and will have difficulty managing a compromise in security.

8.1.6.4 What Are The Users′ Rights and Responsibilities?

The policy should incorporate a statement on the users′ rights and responsibilities concerning the use of the site′s computer systems and services. It should be clearly stated that users are responsible for understanding and respecting the security rules of the systems they are using. The following is a list of topics that you may wish to cover in this area of the policy:

- What guidelines you have regarding resource consumption (whether users are restricted, and if so, what the restrictions are).
- What might constitute abuse in terms of system performance.
- Whether users are permitted to share accounts or let others use their accounts.
- How secret should users keep their passwords.
- How often users should change their passwords and any other password restrictions or requirements.
- Whether you provide backups or expect the users to create their own.
- Disclosure of information that may be proprietary.
- Statement on electronic mail privacy (Electronic Communications Privacy Act).
- Your policy concerning controversial mail or postings to mailing lists or discussion groups (obscenity, harassment, etc.).

- Policy on electronic communications: mail forging, etc.

8.1.6.5 What Happens When the Policy Is Violated?

It is obvious that when any type of official policy is defined, be it related to computer security or not, it will eventually be broken. The violation may occur due to an individual's negligence, accidental mistake, having not been properly informed of the current policy, or not understanding the current policy. It is equally possible that an individual (or group of individuals) may knowingly perform an act that is in direct violation of the defined policy.

When a policy violation has been detected, the immediate course of action should be pre-defined to ensure prompt and proper enforcement. An investigation should be performed to determine how and why the violation occurred. Then the appropriate corrective action should be executed. The type and severity of action taken varies depending on the type of violation that occurred.

8.1.7 Locking In or Out

Whenever a site suffers an incident that compromises computer security, the strategies for reacting may be influenced by two opposing pressures.

If management fears that the site is sufficiently vulnerable, it may choose a protect and proceed strategy. This approach will have as its primary goal the protection and preservation of the site facilities and to provide for normalcy for its users as quickly as possible. Attempts will be made to actively interfere with the intruders processes, prevent further access and begin immediate damage assessment and recovery. This process may involve shutting down the facilities, closing off access to the network, or other drastic measures. The drawback is that unless the intruder is identified directly, they may come back into the site via a different path, or may attack another site.

The alternate approach, pursue and prosecute, adopts the opposite philosophy and goals. The primary goal is to allow intruders to continue their activities at the site until the site can identify the responsible persons. This approach is endorsed by law enforcement agencies and prosecutors. The drawback is that the agencies cannot exempt a site from possible user lawsuits if damage is done to their systems and data.

Prosecution is not the only outcome possible if the intruder is identified. If the culprit is an employee or a student, the organization may choose to take disciplinary actions. The computer security policy needs to spell out the choices and how they will be selected if an intruder is caught.

Careful consideration must be made by site management regarding their approach to this issue before the problem occurs. The strategy adopted might depend upon each circumstance. Or there may be a global policy which mandates one approach in all circumstances. The pros and cons must be examined thoroughly and the users of the facilities must be made aware of the policy so that they understand their vulnerabilities no matter which approach is taken.

The following is a checklists to help a site determine whether or not to adopt protect and proceed.

Protect and Proceed

- If assets are not well protected.
- If continued penetration could result in great financial risk.
- If the possibility or willingness to prosecute is not present.
- If user base is unknown.
- If users are unsophisticated and their work is vulnerable.
- If the site is vulnerable to lawsuits from users.

8.2 Establishing Procedures to Prevent Security Problems

The security policy by itself doesn't say *how* things are protected. The security policy should be a high level document, giving general strategy. The security procedures need to set out, in detail, the precise steps your site will take to protect itself.

The security policy should include a general risk assessment of the types of threats a site is mostly likely to face and the consequences of those threats. Part of doing a risk assessment will include creating a general list of assets that should be protected. This information is critical in devising cost-effective procedures.

It is often tempting to start creating security procedures by deciding on different mechanisms first: our site should have logging on all hosts, call-back modems, and smart cards for all users. This approach could lead to some areas that have too much protection for the risk they face, and other areas that aren't protected enough. Starting with the security policy and the risks it outlines should ensure that the procedures provide the right level of protection for all assets.

8.2.1 Identifing Possible Problems

To determine risk, vulnerabilities must be identified. Part of the purpose of the policy is to aid in shoring up the vulnerabilities and thus decreasing the risk in as many areas as possible.

8.2.1.1 Access Points

Access points are typically used for entry by unauthorized users. Having many access points increases the risk of access to an organization's computer and network facilities. Network links to networks outside the organization allow access into the organization for all others connected to that external network. A network link typically provides access to a large number of network services, and each service has a potential to be compromised. Dialup lines, depending on their configuration, may provide access merely to a login port of a single system. If connected to a terminal server, the dialup line may give access to the entire network. Terminal servers themselves can be a source of problem. Many terminal servers do not require any kind of authentication. Intruders often use terminal servers to disguise their actions, dialing in on a local phone and then using the terminal server to go out to the local network. Some terminal servers are configured so that intruders can telnet in from outside the network, and then telnet back out again, again making it difficult to trace them.

8.2.1.2 Software Bugs

Software will never be bug free. Publicly known security bugs are common methods of unauthorized entry. Part of the solution to this problem is to be aware of the security problems and to update the software when problems are detected. When bugs are found, they should be reported to the vendor so that a solution to the problem can be implemented and distributed.

8.2.1.3 Insider Threats

An insider to the organization may be a considerable threat to the security of the computer systems. Insiders often have direct access to the computer and network hardware components. The ability to access the components of a system makes most systems easier to compromise. Most desktop workstations can be easily manipulated so that they grant privileged access. Access to a local area network provides the ability to view possibly sensitive data traversing the network.

8.2.2 Choose Controls to Protect Assets in a Cost-Effective Way

After establishing what is to be protected, and assessing the risks these assets face, it is necessary to decide how to implement the controls which protect these assets. The controls and protection mechanisms should be selected in a way so as to adequately counter the threats found during risk assessment, and to implement those controls in a cost-effective manner. It makes little sense to spend an exorbitant sum of money and overly constrict the user base if the risk of exposure is very small.

8.2.2.1 Choose the Right Set of Controls

The controls that are selected represent the physical embodiment of your security policy. They are the first and primary line of defense in the protection of your assets. It is therefore most important to ensure that the controls that you select are the right set of controls. If the major threat to your system is outside penetrators, it probably doesn′t make much sense to use biometric devices to authenticate your regular system users. On the other hand, if the major threat is unauthorized use of computing resources by regular system users, you′ll probably want to establish very rigorous automated accounting procedures.

8.2.2.2 Use Common Sense

Common sense is the most appropriate tool that can be used to establish your security policy. Elaborate security schemes and mechanisms are impressive, and they do have their place, yet there is little point in investing money and time on an elaborate implementation scheme if the simple controls are forgotten. For example, no matter how elaborate a system you put into place on top of existing security controls, a single user with a poor password can still leave your system open to attack.

8.2.2.3 Use Multiple Strategies to Protect Assets

Another method of protecting assets is to use multiple strategies. In this way, if one strategy fails or is circumvented, another strategy comes into play to continue protecting the asset. By using several simpler strategies, a system can often be made more secure than if one very sophisticated method were used in its place. For example, dial-back modems can be used in conjunction with traditional logon mechanisms. Many similar approaches could be devised that provide several levels of protection for assets.

However, it's very easy to go overboard with extra mechanisms. One must keep in mind exactly what it is that needs to be protected.

8.3 Physical Security

It is a given in computer security that if the system itself is not physically secure, nothing else about the system can be considered secure. With physical access to a machine, an intruder can halt the machine, bring it back up in privileged mode, replace or alter the disk , plant virus programs, or take any number of other undesirable (and hard to prevent) actions. Critical communications links, important servers, and other key machines should be located in physically secure areas. Some security systems (such as Kerberos) require that the machine be physically secure. If you cannot physically secure machines, care should be taken about trusting those machines. Sites should consider limiting access from non-secure machines to more secure machines. In particular, allowing trusted access from these kinds of hosts is particularly risky. For machines that seem or are intended to be physically secure, care should be taken about who has access to the machines. Remember that custodial and maintenance staff often have keys to rooms and may not knowingly allow access to unauthorized individuals.

8.3.1 Procedures to Recognize Unauthorized Activity

Several simple procedures can be used to detect most unauthorized uses of a computer system. These procedures use tools provided with the operating system by the vendor, or tools publicly available from other sources.

8.3.1.1 Monitoring System Use

System monitoring can be done either by a system administrator or by software written for the purpose. Monitoring a system involves looking at several parts of the system and searching for anything unusual. The most important thing about monitoring system use is that it be done on a regular basis. Picking one day out of the month to monitor the system is pointless, since a security breach can be isolated to a matter of hours. Only by maintaining a constant vigil can you expect to detect security violations in time to react to them.

8.3.2 Tools for Monitoring the System

8.3.2.1 Logging

Most operating systems store numerous bits of information in log files. Examination of these log files on a regular basis is often the first line of defense in detecting unauthorized use of the system.

Compare lists of currently logged in users and past login histories: Most users typically log in and out at roughly the same time each day. An account logged in outside the ″normal″ time for the account may be in use by an intruder.

Many systems maintain accounting records for billing purposes: These records can also be used to determine usage patterns for the system; unusual accounting records may indicate unauthorized use of the system.

System logging facilities, such as the UNIX syslog: Utility should be checked for unusual error messages from system software. For example, a large number of failed login attempts in a short period of time may indicate someone trying to guess passwords.

Operating system commands: That list currently executing processes can be used to detect users running programs they are not authorized to use, as well as to detect unauthorized programs that have been started by an intruder.

8.3.2.2 Monitoring Software

Other monitoring tools can easily be constructed using standard operating system software, by using several, often unrelated, programs together. For example, checklists of file ownerships and permission settings can be constructed (for example, with ls and find on UNIX) and stored offline. These lists can then be reconstructed periodically and compared against the master checklist (on UNIX, by using the diff utility). Differences may indicate that unauthorized modifications have been made to the system.

8.3.2.3 Other Tools

Other tools can also be used to monitor systems for security violations, although this is not their primary purpose. For example, network monitors can be used to detect and log connections from unknown sites.

8.3.3 Vary the Monitoring Schedule

The task of system monitoring is not as daunting as it may seem. System administrators can execute many of the commands used for monitoring periodically throughout the day during idle moments (for example while talking on the telephone), rather than spending fixed periods of each day monitoring the system. By executing the commands frequently, you will rapidly become used to seeing normal output, and will easily spot things that are out of the ordinary. In addition, by running various monitoring commands at different times throughout the day, you make it hard for an intruder to predict your actions. For example, if an intruder knows that each day at 5:00 p.m. the system is checked to see that everyone has logged off, he will simply wait until after the check has completed before logging in. But the intruder cannot guess when a system administrator might type a command to display all logged in users, and thus he runs a much greater risk of detection.

Despite the advantages that regular system monitoring provides, some intruders will be aware of the standard logging mechanisms in use on systems they are attacking. They will actively pursue and attempt to disable monitoring mechanisms. Regular monitoring therefore is useful in detecting intruders, but does not provide any guarantee that your system is secure. Also, monitoring should not be considered an *infallible* method of detecting unauthorized use.

8.3.3.1 Define Actions to Take when Unauthorized Activity Is Supected

The procedures for dealing with these types of problems should be written down. Who has authority to decide what actions will be taken? Should law enforcement be involved? Should your organization cooperate with other sites in trying to track down an intruder? Whether you decide to lock out or pursue intruders, you should have tools and procedures ready to apply. It is best to work up these tools and procedures before you need them. Don't wait until an intruder is on your system to figure out how to track the intruder's actions; you will be busy enough if an intruder strikes.

8.3.4 Communicating Security Policy

Security policies, in order to be effective, must be communicated to both the users of the system and the system maintainers.

8.3.4.1 Educating the Users

Users should be made aware of how the computer systems are expected to be used, and how to protect themselves from unauthorized users.

Proper Account/Workstation Use: All users should be informed about what is considered the "proper" use of their account or workstation. This can most easily be done at the time a user receives their account by giving them a policy statement. Proper use policies typically dictate things such as whether or not the account or workstation may be used for personal activities (such as checkbook balancing or letter writing), whether profit-making activities are allowed, whether game playing is permitted, and so on. These policy statements may also be used to summarize how the computer facility is licensed and what software licenses are held by the institution; for example, many universities have educational licenses which explicitly prohibit commercial uses of the system.

Account/Workstation Management Procedures: Each user should be told how to properly manage their account and workstation. This includes explaining how to protect files stored on the system, how to log out or lock the terminal or workstation, and so on. Much of this information is typically covered in the beginning user documentation provided by the operating system vendor, although many sites elect to supplement this material with local information. If your site offers dial-up modem access to the computer systems, special care must be taken to inform users of the security problems inherent in providing this access. Issues such as making sure to log out before hanging up the modem should be covered when the user is initially given dial-up access. Likewise, access to the systems via local and wide area networks presents its own set of security problems which users should be made aware of. Files that grant trusted host or trusted user status to remote systems and users should be carefully explained.

Determining Account Misuse: Users should be told how to detect unauthorized access to their account. If the system prints the last login time when a user logs in, he or she should be told to check that time and note whether or not it agrees with the last time he or she actually logged in. Command interpreters on some systems maintain histories of the last several commands executed. Users should check these histories to be sure someone has not executed other commands with their account.

Problem Reporting Procedures: A procedure should be developed to enable users to report suspected misuse of their accounts or other misuse they may have noticed. This can be done either by providing the name and telephone number of a system administrator who manages security of the computer system, or by creating an electronic mail address to which users can address their problems.

8.3.4.2 Educating the Host Administrators

In many organizations, computer systems are administered by a wide variety of people. These administrators must know how to protect their own systems from attack and unauthorized use, as well as how to communicate successful penetration of their systems to other administrators as a warning.

Account Management Procedures: Care must be taken when installing accounts on the system in order to make them secure. When installing a system from distribution media, the password file should be examined for standard accounts provided by the vendor. Many vendors provide accounts for use by system services or field service personnel. These accounts typically have either no password or one which is common knowledge. These accounts should be given new passwords if they are needed, or disabled or deleted from the system if they are not. Accounts without passwords are generally very dangerous since they allow anyone to access the system.

Even accounts that do not execute a command interpreter (accounts that exist only to see who is logged in to the system) can be compromised if set up incorrectly. A related concept is that of anonymous file transfer (FTP), which allows workstations users from all over the network to access your system to retrieve files from (usually) a protected disk area. You should carefully weigh the benefits that an account without a password provides against the security risks of providing such access to your system. If the operating system provides a shadow password facility that stores passwords in a separate file accessible only to privileged users, this facility should be used. It protects passwords by hiding their encrypted values from unprivileged users. This prevents an attacker from copying your password file to his or her machine and then attempting to break the passwords at his or her leisure. Keep track of who has access to privileged user accounts (the root user ID on UNIX or the MAINT user ID on VMS). Whenever a privileged user leaves the organization or no longer has need of the privileged account, the passwords on all privileged accounts should be changed.

Configuration Management Procedures: When installing a system from the distribution media or when installing third-party software, it is important to check the installation carefully. Many installation procedures assume a trusted site, and hence will install files with world-writeable permission enabled, or otherwise compromise the security of files. Network services should also be examined carefully when first installed. Many vendors provide default network permission files which imply that all outside hosts are to be trusted, which is rarely the case when connected to wide area networks such as the Internet.

Many intruders collect information on the vulnerabilities of particular system versions. The older a system, the more likely it is that there are security problems in that version that have since been fixed by the vendor in a later

release. For this reason, it is important to weigh the risks of not upgrading to a new operating system release (thus leaving security holes unplugged) against the cost of upgrading to the new software (possibly breaking third-party software, etc.).

Bug fixes from the vendor should be weighed in a similar fashion, with the added note that security fixes from a vendor usually address fairly serious security problems. Other bug fixes, received via network mailing lists and the like, should usually be installed, but not without careful examination. Never install a bug fix unless you're sure you know what the consequences of the fix are; there's always the possibility that an intruder has suggested a fix which actually gives him or her access to your system.

Recovery Procedures - Backups: It is impossible to overemphasize the need for a good backup strategy. File system backups not only protect you in the event of hardware failure or accidental deletions, but they also protect you against unauthorized changes made by an intruder. Without a copy of your data the way it's supposed to be, it can be difficult to undo something an attacker has done. Backups, especially if run daily, can also be useful in providing a history of an intruder's activities. Looking through old backups can establish when your system was first penetrated. Intruders may leave files around which, although deleted later, are captured on the backup tapes. Backups can also be used to document an intruder's activities to law enforcement agencies if necessary. A good backup strategy will dump the entire system to tape at least once a month. Partial (or incremental) dumps should be done at least twice a week, and ideally they should be done daily. Commands specifically designed for performing file system backups (UNIX dump or VMS BACKUP command) should be used in preference to other file copying commands, since these tools are designed with the express intent of restoring a system to a known state.

8.3.4.3 Problem Reporting Procedures

As with users, system administrators should have a defined procedure for reporting security problems. In large installations, this is often done by creating an electronic mail alias that contains the names of all system administrators in the organization. Other methods include setting up some sort of response team similar to the CERT, or establishing a hotline serviced by an existing support group.

8.3.5 Resources to Prevent Security Breaches

These are some of the resources to prevent security breaches.

8.3.5.1 Network Connections

Some sites will be connected only to other sites within the same organization and will not have the ability to connect to other networks. Sites such as these are less susceptible to threats from outside their own organization, although intrusions may still occur via paths such as dial-up modems. On the other hand, many other organizations will be connected to other sites via much larger networks, such as the Internet. These sites are susceptible to the entire range of threats associated with a networked environment. The risks of connecting to outside networks must be weighed against the benefits. It may be desirable to limit connection to outside networks to those hosts which do not store sensitive material, keeping vital machines (such as those which maintain company payroll or inventory

systems) isolated. If there is a need to participate in a wide area network (WAN), consider restricting all access to your local network through a single system.

8.3.5.2 Firewalls

A firewall is a system or group of systems that enforces an access control policy between two networks. The actual means by which this is accomplished varies widely, but in principle, the firewall can be thought of as a pair of mechanisms: one which exists to block traffic, and the other which exists to permit traffic. Some firewalls place a greater emphasis on blocking traffic, while others emphasize permitting traffic. Probably the most important thing to recognize about a firewall is that it implements an access control policy.

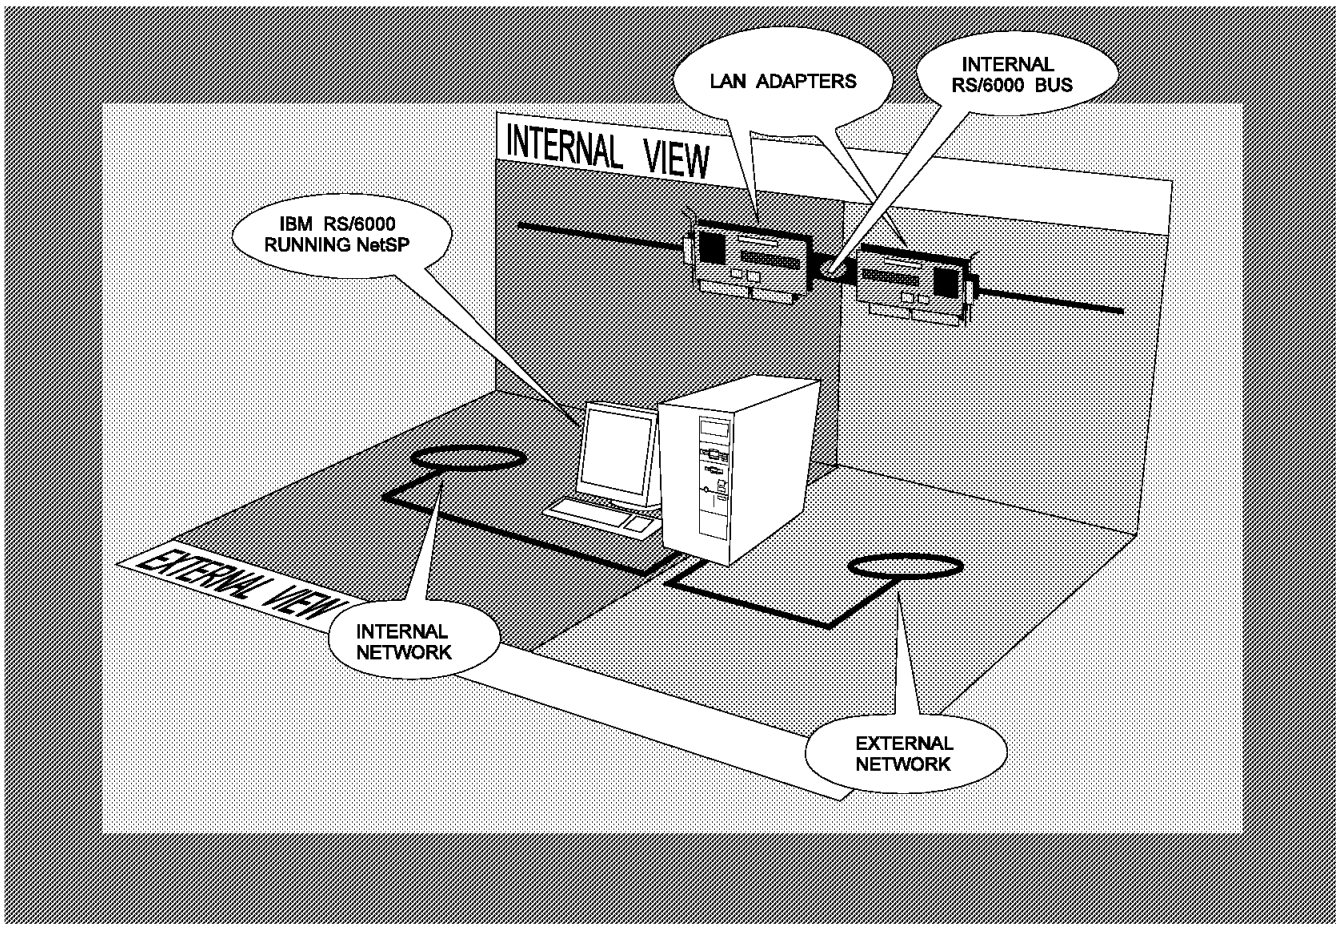


Figure 166. Firewall. This figure shows the IBM NetSP firewall solution, running on the RS/6000 platform.

The Internet, like any other society, is plagued with the kind of jerks who enjoy the electronic equivalent of writing on other people′s walls with spraypaint, tearing mailboxes off, or just sitting in the street blowing their car horns. Some people try to get real work done over the Internet, and others have sensitive or proprietary data they must protect. Usually, a firewall′s purpose is to keep the jerks out of your network while still letting you get your job done.

The firewall can act as your corporate ambassador to the Internet. Many corporations use their firewall systems as a place to store public information

about corporate products and services, files to download, bug fixes, and so forth.

Several of these systems have become important parts of the Internet service structure (UUnet.uu.net, whitehouse.gov, gatekeeper.dec.com) and have reflected well on their organizational sponsors.

Some firewalls permit only e-mail traffic through them, thereby protecting the network against any attacks other than attacks against the e-mail service. Other firewalls provide less strict protections, and block services that are known to be problems. Generally, firewalls are configured to protect against unauthenticated interactive logins from the outside world. This, more than anything, helps prevent vandals from logging into machines on your network.

More elaborate firewalls block traffic from the outside to the inside, but permit users on the inside to communicate freely with the outside. The firewall can protect you against any type of network-borne attack if you unplug it. Firewalls are also important since they can provide a single choke point where security and audit can be imposed. Unlike in a situation where a computer system is being attacked by someone dialing in with a modem, the firewall can act as an effective phone tap and tracing tool.

Firewalls provide an important logging and auditing function; often they provide summaries to the administrator about what kinds and amount of traffic passed through it, how many attempts there were to break into it, etc. Firewalls can't protect against attacks that don't go through the firewall. Many corporations that connect to the Internet are very concerned about proprietary data leaking out of the company through that route. Unfortunately for those concerned, a magnetic tape can just as effectively be used to export data. Many organizations that are terrified (at a management level) of Internet connections have no coherent policy about how dial-in access via modems should be protected. It's silly to build a 6-foot thick steel door when you live in a wooden house, but there are a lot of organizations out there buying expensive firewalls and neglecting the numerous other back-doors into their network.

For a firewall to work, it must be a part of a consistent overall organizational security architecture. Firewall policies must be realistic, and reflect the level of security in the entire network. For example, a site with top secret or classified data doesn't need a firewall at all: they shouldn't be hooking up to the Internet in the first place, or the systems with the secret data should be isolated from the rest of the corporate network. Another thing a firewall can't really protect you against is traitors to your network. While industrial spies might export information through your firewall, they are just as likely to export it through a telephone, fax machine, or floppy disk. Floppy disks are a far more likely means for information to leak from your organization than a firewall.

Firewalls also cannot protect you against stupidity. Users who reveal sensitive information over the telephone are good targets for social engineering; attackers may be able to break into your network by completely bypassing your firewall, if they can find a helpful employee inside who can be fooled into giving access to a modem pool. Conceptually, there are two types of firewalls:

- Network Level

- Application Level

They are not as different as you might think, and the latest technologies are blurring the distinction to the point where it's no longer clear if either one is better than the other. As always, you need to be careful to pick the type that meets your needs.

Network level firewalls: Network level firewalls generally make their decisions based on the source, destination addresses and ports in individual IP packets. A simple router is the traditional network level firewall, since it is not able to make particularly sophisticated decisions about what a packet is actually talking to or where it actually came from. Modern network level firewalls have become increasingly sophisticated, and now maintain internal information about the state of connections passing through them, the contents of some of the data streams, and so on. One thing that's an important distinction about many network level firewalls is that they route traffic directly though them, so to use one you usually need to have a validly assigned IP address block. Network level firewalls tend to be very fast and tend to be very transparent to users.

Application level firewalls: Application level firewalls generally are hosts running proxy servers, which permit no traffic directly between networks, and which perform elaborate logging and auditing of traffic passing through them. Since the proxy applications are software components running on the firewall, it is a good place to do lots of logging and access control. Application level firewalls can be used as network address translators, since traffic goes in one side and out the other, after having passed through an application that effectively masks the origin of the initiating connection. Having an application in the way in some cases may impact performance and may make the firewall less transparent. Early application level firewalls such as those built using the TIS firewall toolkit, are not particularly transparent to end users and may require some training. Modern application level firewalls are often fully transparent. Application level firewalls tend to provide more detailed audit reports and tend to enforce more conservative security models than network level firewalls.

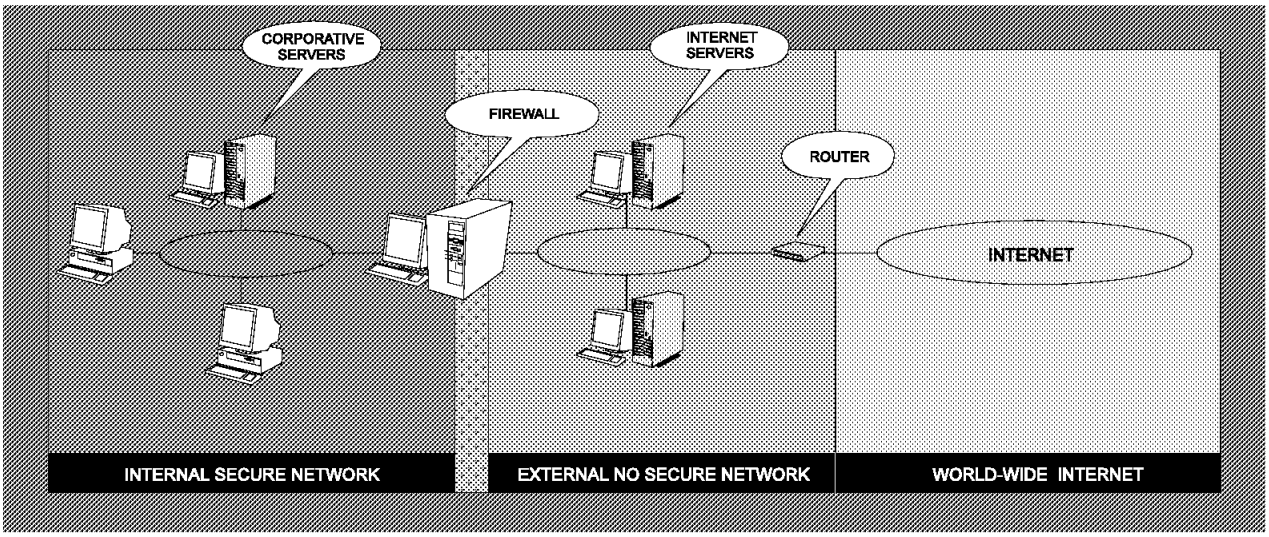


Figure 167. Firewall Solution. This figure shows a typical corporative secure network

Proxy servers: A proxy server is an application that mediates traffic between a protected network and the Internet. Proxies are often used instead of router-based traffic controls, to prevent traffic from passing directly between networks. Many proxies contain extra logging or support for user authentication. Since proxies must understand the application protocol being used, they can also implement protocol-specific security. An FTP proxy might be configurable to permit incoming FTP and block outgoing FTP. Proxy servers are application specific. In order to support a new protocol via a proxy, a proxy must be developed for it.

Socks servers: A socks is a generic proxy system that can be compiled into a client-side application to make it work through a firewall. Its advantage is that it′s easy to use, but it doesn′t support the addition of authentication hooks or protocol specific logging. For more information on socks, see ftp.nec.com:/pub/security/socks.cstc.

Using a firewall with DNS systems: Some organizations want to hide DNS names from the outside. Many experts don′t think hiding DNS names is worthwhile, but if site/corporate policy mandates hiding domain names, this is one approach that is known to work. Another reason you may have to hide domain names is if you have a non-standard addressing scheme on your internal network. In that case, you have no choice but to hide those addresses. Don′t fool yourself into thinking that if your DNS names are hidden that it will slow attackers down if they break into your firewall. Information about what is on your network is too easily gleaned from the networking layer itself. If you want an interesting demonstration of this, ping the subnet broadcast address on your LAN and then type `arp -a`. Note also that hiding names in the DNS doesn′t address the problem of host names leaking out in mail headers, news articles, etc. This approach is one of many, and is useful for organizations that wish to hide their host names from the Internet. The success of this approach lies on the fact that DNS clients on a machine don′t have to talk to a DNS server on that same machine. In other words, just because there′s a DNS server on a machine, there′s nothing wrong with (and there are often advantages to) redirecting that machine′s DNS client activity to a DNS server on another machine.

First, you set up a DNS server on the bastion host that the outside world can talk to. You set this server up so that it claims to be authoritative for your domains. In fact, all this server knows is what you want the outside world to know; the names and addresses of your gateways, your wildcard MX records, and so forth. This is the public server.

Then, you set up a DNS server on an internal machine. This server also claims to be authoritiative for your domains; but unlike the public server, this one is telling the truth. This is your normal nameserver, into which you put all your normal DNS stuff. You also set this server up to forward queries that it can′t resolve to the public server.

Finally, you set up all your DNS clients, including the ones on the machine with the public server, to use the internal server. An internal client asking about an internal host asks the internal server, and gets an answer; an internal client asking about an external host asks the internal server, which asks the public server, which asks the Internet, and the answer is relayed back. A client on the public server works just the same way. An external

client, however, asking about an internal host gets back the restricted answer from the public server.

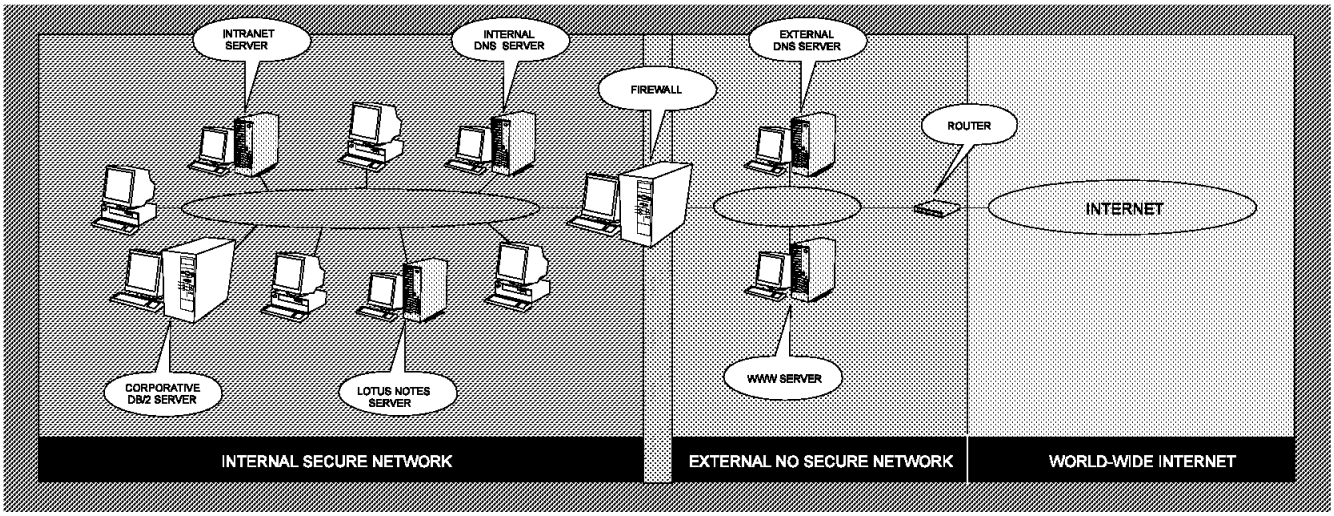


Figure 168. Hidden DNS names. This figure shows a corporative Internet solution using internal and external DNS server

This approach assumes that there's a packet filtering firewall between these two servers that will allow them to talk DNS to each other, but otherwise restricts DNS between other hosts. Another trick that's useful in this scheme is to employ wildcard PTR records in your IN-ADDR.ARPA domains. These cause an an address-to-name lookup for any of your nonpublic hosts to return something such as "unknown.YOUR.DOMAIN" rather than an error. This satisfies anonymous FTP sites like ftp.uu.net that insist on having a name for the machines they talk to. This may fail when talking to sites that do a DNS cross-check in which the host name is matched against its address and vice versa.

Using FTP through the firewall: Generally, making FTP work through the firewall is done either using a proxy server such as the firewall toolkit's ftp-gw or by permitting incoming connections to the network at a restricted port range, and otherwise restricting incoming connections using something such as established screening rules. The FTP client is then modified to bind the data port to a port within that range. This entails being able to modify the FTP client application on internal hosts. In some cases, if FTP downloads are all you wish to support, you might want to consider declaring FTP a dead protocol and letting you users download files via the Web instead. The user interface certainly is nicer, and it gets around the ugly callback port problem. If you choose the FTP-via-Web approach, your users will be unable to FTP files out, which, depending on what you are trying to accomplish, may be a problem.

Using Telnet through the firewall: Telnet is generally supported either by using an application proxy such as the firewall toolkit's tn-gw, or by simply configuring a router to permit outgoing connections using something like the established screening rules. Application proxies could be in the form of a stand-alone proxy running on the bastion host, or in the form of a SOCKS server and a modified client.

Using GOPHER, ARCHIE and other services through the firewall: The majority of firewall administrators choose to support GOPHER and ARCHIE through Web proxies, instead of directly. The Web′s tendency to make everything on the Internet look like a Web service is both a blessing and a curse. There are many new services constantly cropping up. Often they are misdesigned or are not designed with security in mind, and their designers will cheerfully tell you if you want to use them you need to let port xxx through your router. Unfortunately, not everyone can do that, and so a number of interesting new toys are difficult to use for people behind firewalls. Things like RealAudio, which require direct UDP access, are particularly egregious examples. The thing to bear in mind if you find yourself faced with one of these problems is to find out as much as you can about the security risks that the service may present, before you just allow it through. It′s quite possible the service has no security implications. It′s equally possible that it has undiscovered holes you could drive a truck through.

Using X-WINDOWS through the firewall: X-WINDOWS is a very useful system, but unfortunately it has some major security flaws. Remote systems that can gain or spoof access to a workstation′s X display can monitor keystrokes that a user enters, download copies of the contents of their windows, etc. While attempts have been made to overcome them, it is still entirely too easy for an attacker to interfere with a user′s X display. Most firewalls block all X traffic.

Sourced routed traffic: Normally, the route a packet takes from its source to its destination is determined by the routers between the source and destination. The packet itself only says where it wants to go (the destination address), and nothing about how it expects to get there. There is an optional way for the sender of a packet (the source) to include information in the packet that tells the route the packet should get to its destination; thus the name source routing. For a firewall, source routing is noteworthy, since an attacker can generate traffic claiming to be from a system inside the firewall. In general, such traffic wouldn′t route to the firewall properly, but with the source routing option, all the routers between the attacker′s machine and the target will return traffic along the reverse path of the source route. Implementing such an attack is quite easy; so firewall builders should not discount it as unlikely to happen. In practice, source routing is used very little. In fact, generally the main legitimate use is in debugging network problems or routing traffic over specific links for congestion control for specialized situations. When building a firewall, source routing should be blocked at some point. Most commercial routers incorporate the ability to block source routing specifically, and many versions of UNIX that might be used to build firewall bastion hosts have the ability to disable or ignore source routed traffic

Denial of service: Denial of service is when someone decides to make your network or firewall useless by disrupting it, crashing it, jamming it, or flooding it. The problem with denial of service on the Internet is that it is impossible to prevent. The reason has to do with the distributed nature of the network; every network node is connected via other networks which in turn connect to other networks, etc. A firewall administrator or ISP only has control of a few of the local elements within reach. An attacker can always disrupt a connection upstream from where the victim controls it. In other words, if someone wanted to take a network off the air, they could do it

either by taking the network off the air, or by taking the networks it connects to off the air, etc. There are many, many, ways someone can deny service, ranging from the complex to the brute-force. If you are considering using Internet for a service which is absolutely time or mission-critical, you should consider your fall-back position in the event that the network is down or damaged.

8.3.5.3 IBM Secure Network Gateway

The IBM Internet Connection Secured Network Gateway (SNG) is based on research at IBM′s Yorktown Research Laboratory and experience running large networks for more than eight years. SNG support includes:

- Secure IP tunnels
- IP filters
- Proxy servers
- Socks servers
- Secured services, such as the Domain Name Service or mail handling

Secure IP tunnels use an encapsulation scheme to insert IP packets and their headers into encrypted IP packets. IP tunnels let administrators set security policy without requiring users to get involved. With IP tunnels, the firewall at the sending end of the tunnel encloses the sender′s information into encrypted packets and sends the packets to the receiving firewall. The receiving firewall removes the encapsulation. The path between firewalls forms a secure tunnel through the Internet. The firewall administrators determine the levels of protection and the types of information protected at the IP address and port level. Obviously, the ends of the tunnel have to agree, or the packets will be unintelligible and discarded. Secure IP tunnels are an effective way to implement a security policy between a reasonable number of homogenous firewalls.

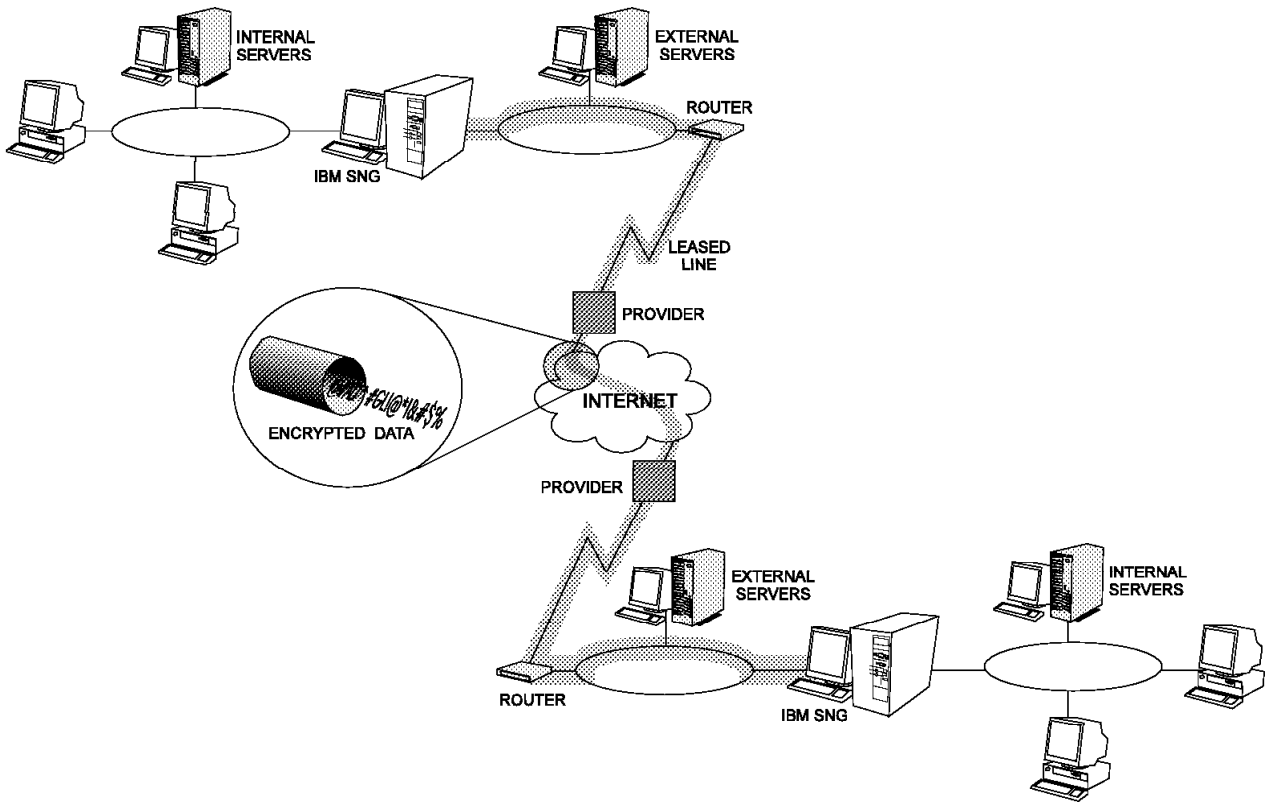


Figure 169. IBM Secure Network Gateway. This figure show the tunning feature included on IBM SNG

A secure network gateway (SNG) limits private network users access to the public network with a command shell that restricts commands like Telnet, Mosaic, and Gopher. SNG does not include any commands that let the user look at or modify the firewall. The advantage of the proxy server is that users do not have to have any special client code. They use the same code they would use in a non-proxy implementation. However, each application requires a double connection: one to the proxy, and one to the Ultimate destination. This can be time consuming, and has a performance impact. Running the Domain Name Server on the SNG firewall hides private network hosts from the nonsecure world and prevents name resolution requests from flowing across the gateway uncontrolled. SNG also provides a simplified sendmail daemon that acts as a mail relay. When administrators define an SNG Domain Name Service, they can also specify a secure network mail gateway. Only the SNG mail server is advertised outside the private network. The SNG mail gateway can forward mail to a standard mail gateway within the private network, providing the best of both worlds: full-function mail services within the private network with a secure mail interface to the public network.

Principal features about the IBM Network Secure Network Gateway:

- **Alarm facilities:** The IBM firewall allows you to actively monitor security events at the firewall and generate real-time notification to the network administrator.

- **Advanced filtering capability:** Filters are used to control packet flows based upon criteria, such as IP source or destination address range, TCP ports, UDP, ICMP, and TCP responses. Filters are transparent to users, and are a powerful way to deny access to specific locations within your network.
- **Application gateway proxy:** Using either Telnet or FTP, users can access the IBM firewall, where their identity is authenticated. After verifying a user′s identity, the firewall allows the user to launch any TCP/IP application that the user is authorized to access, such as FTP, Gopher, and WHOIS. All packets flowing from the IBM firewall carry the IP address of the firewall as the originating address. So, the gateway proxy server hides the IP addresses of your internal network from the outside world. It also allows you to grant trust on the basis of individual users, rather than on the basis of an IP address.
- **SOCKS server:** Applications, running on hosts and workstations within your secured network that use the SOCKS API, can use the SOCKS server on the IBM firewall. SOCKS can be used to provide a transparent means of controlling access to the Internet, while, at the same time, hiding the IP addresses of your internal network from the outside world.
- **Domain Name Server:** The external Domain Name Server presents your corporate domain name to the Internet. The outside world can′t see the structure of your network or the names and addresses of your internal hosts.
- **Mail service:** The IBM firewall supports forwarding of authorized Simple Mail Transfer Protocol (SMTP) e-mail to an e-mail server in the secure network.
- **Strong authentication:** The IBM firewall offers various methods for authenticating clients. You can use a password or more sophisticated methods, like Digital Pathways′ SecureNet card or Security Dynamics′ SecurID card.
- **Services and support:** IBM offers expert professional services to properly set up a secure firewall platform, write the permit-or-deny rules that reflect your company′s security policy, and train your operations staff to administer the firewall. IBM also offers a complement of support line services to help keep your IBM firewall maintained.

IBM SNG hardware requirements

- RISC System/6000 supported by AIX/6000 Version 3.2.5 or 4.1.3, operating systems with 1GB disk space and at least 32 MB of memory
- At least two communication hardware adapters supported by the TCP/IP protocol stack
- 6 MB available for programs

IBM SNG software requirements

- AIX/6000 Version 3.2.5 or 4.1.3

8.3.5.4 Glossary of Firewall-Related Terms

Abuse of privilege: When a user performs an action that they should not have according to organizational policy or law.

Application-level firewall: A firewall system in which service is provided by processes that maintain complete TCP connection state and sequencing. Application level firewalls often re-address traffic so that outgoing traffic appears to have originated from the firewall, rather than the internal host.

Authentication: The process of determining the identity of a user that is attempting to access a system.

Authentication token: A portable device used for authenticating a user. Authentication tokens operate by challenge/response, time-based code sequences, or other techniques. This may include paper-based lists of one-time passwords.

Authorization: The process of determining what types of activities are permitted. Usually, authorization is in the context of authentication: once you have authenticated a user, they may be authorized different types of access or activity.

Bastion host: A system that has been hardened to resist attack, and which is installed on a network in such a way that it is expected to potentially come under attack. Bastion hosts are often components of firewalls, or may be outside Web servers or public access systems. Generally, a bastion host is running some form of general purpose operating system (for examplw, UNIX, VMS, WNT, etc.) rather than a ROM-based or firmware operating system.

Challenge/response: An authentication technique whereby a server sends an unpredictable challenge to the user, who computes a response using some form of authentication token.

Chroot: A technique under UNIX whereby a process is permanently restricted to an isolated subset of the file system.

Cryptographic checksum: A one-way function applied to a file to produce a unique ″fingerprint″ of the file for later reference. Checksum systems are a primary means of detecting file system tampering on UNIX.

Data Driven attack: A form of attack in which the attack is encoded in innocuous-seeming data which is executed by a user or other software to implement an attack. In the case of firewalls, a data-driven attack is a concern since it may get through the firewall in data form and launch an attack against a system behind the firewall.

Defense in depth: The security approach whereby each system on the network is secured to the greatest possible degree. May be used in conjunction with firewalls.

DNS spoofing: Assuming the DNS name of another system by either corrupting the name service cache of a victim system, or by compromising a domain name server for a valid domain.

Dual homed gateway: A dual homed gateway is a system that has two or more network interfaces, each of which is connected to a different network.

In firewall configurations, a dual homed gateway usually acts to block or filter some or all of the traffic trying to pass between the networks.

Encrypting Router: see tunneling router and virtual network perimeter.

Firewall: A system or combination of systems that enforces a boundary between two or more networks.

Host-based security: The technique of securing an individual system from attack. Host-based security is operating system and version dependent.

Insider attack: An attack originating from inside a protected network.

Intrusion detection: Detection of break-ins or break-in attempts either manually or via software expert systems that operate on logs or other information available on the network.

IP Spoofing: An attack whereby a system attempts to illicitly impersonate another system by using its IP network address.

IP Splicing / hijacking: An attack whereby an active, established, session is intercepted and co-opted by the attacker. IP splicing attacks may occur after an authentication has been made, permitting the attacker to assume the role of an already authorized user. Primary protections against IP splicing rely on encryption at the session or network layer.

Least privilege: Designing operational aspects of a system to operate with a minimum amount of system privilege. This reduces the authorization level at which various actions are performed and decreases the chance that a process or user with high privileges may be caused to perform unauthorized activity resulting in a security breach.

Logging: The process of storing information about events that occurred on the firewall or network.

Log retention: How long audit logs are retained and maintained.

Log processing: How audit logs are processed, searched for key events, or summarized.

Network-level firewall: A firewall in which traffic is examined at the network protocol packet level.

Perimeter-based security: The technique of securing a network by controlling access to all entry and exit points of the network.

Policy: Organization-level rules governing acceptable use of computing resources, security practices, and operational procedures.

Proxy: A software agent that acts on behalf of a user. Typical proxies accept a connection from a user, make a decision as to whether or not the user or client IP address is permitted to use the proxy, perhaps does additional authentication, and then completes a connection on behalf of the user to a remote destination.

Screened host: A host on a network behind a screening router. The degree to which a screened host may be accessed depends on the screening rules in the router.

Screened subnet: A subnet behind a screening router. The degree to which the subnet may be accessed depends on the screening rules in the router.

Screening router: A router configured to permit or deny traffic based on a set of permission rules installed by the administrator.

Session stealing: See IP splicing.

Trojan horse: A software entity that appears to do something normal but which, in fact, contains a trap door or attack program.

Tunneling Router: A router or system capable of routing traffic by encrypting it and encapsulating it for transmission across an untrusted network for eventual de-encapsulation and decryption.

Social engineering: An attack based on deceiving users or administrators at the target site. Social engineering attacks are typically carried out by telephoning users or operators and pretending to be an authorized user, to attempt to gain illicit access to systems.

Virtual network perimeter: A network that appears to be a single protected network behind firewalls, which actually encompasses encrypted virtual links over untrusted networks.

Virus: A self-replicating code segment. Viruses may or may not contain attack programs or trap doors.

8.3.5.5 Confidentiality

Confidentiality, the act of keeping things hidden or secret, is one of the primary goals of computer security practitioners. Several mechanisms are provided by most modern operating systems to enable users to control the dissemination of information. Depending upon where you work, you may have a site where everything is protected, or a site where all information is usually regarded as public, or something in-between. Most sites lean toward the in-between, at least until some penetration has occurred. Generally, there are three instances in which information is vulnerable to disclosure: when the information is stored on a computer system, when the information is in transit to another system (on the network), and when the information is stored on backup tapes. The first of these cases is controlled by file permissions, access control lists, and other similar mechanisms. The last can be controlled by restricting access to the backup tapes (by locking them in a safe, for example). All three cases can be helped by using encryption mechanisms.

8.3.5.6 Encryption (Hardware and Software)

Encryption is the process of taking information that exists in some readable form and converting it into a non-readable form. There are several types of commercially available encryption packages in both hardware and software forms. Hardware encryption engines have the advantage that they are much faster than the software equivalent, yet because they are faster, they are of greater potential benefit to an attacker who wants to execute a brute-force attack on your encrypted information. The advantage of using encryption is

that, even if other access control mechanisms (passwords, file permissions, etc.) are compromised by an intruder, the data is still unusable.

Naturally, encryption keys and the like should be protected at least as well as account passwords. Information in transit (over a network) may be vulnerable to interception as well. Several solutions to this exist, ranging from simply encrypting files before transferring them (end-to-end encryption) to special network hardware which encrypts everything it sends without user intervention (secure links). The Internet as a whole does not use secure links, thus end-to-end encryption must be used if encryption is desired across the Internet.

Data Encryption Standard (DES): DES is perhaps the most widely used data encryption mechanism today. Many hardware and software implementations exist, and some commercial computers are provided with a software version. DES transforms plain text information into encrypted data (or ciphertext) by means of a special algorithm and seed value called a key. So long as the key is retained (or remembered) by the original user, the ciphertext can be restored to the original plain text. One of the pitfalls of all encryption systems is the need to remember the key under which a thing was encrypted (this is not unlike the password problem discussed elsewhere in this document). If the key is written down, it becomes less secure. If forgotten, there is little (if any) hope of recovering the original data. Most UNIX systems provide a DES command that enables a user to encrypt data using the DES algorithm.

Crypt: Similar to the DES command, the UNIX crypt command allows a user to encrypt data. Unfortunately, the algorithm used by crypt is very insecure (based on the World War II Enigma device), and files encrypted with this command can be decrypted easily in a matter of a few hours. Generally, use of the crypt command should be avoided for any but the most trivial encryption tasks.

Privacy Enhanced Mail: Electronic mail normally transits the network in the clear (anyone can read it). This is obviously not the optimal solution. Privacy enhanced mail provides a means to automatically encrypt electronic mail messages so that a person eavesdropping at a mail distribution node is not (easily) capable of reading them. Several privacy enhanced mail packages are currently being developed and deployed on the Internet. The Internet Activities Board Privacy Task Force has defined a draft standard, elective protocol for use in implementing privacy enhanced mail.

8.3.5.7 Origin Authentication

We mostly take it on faith that the header of an electronic mail message truly indicates the originator of a message. However, it is easy to forge the source of a mail message. Origin authentication provides a means to be certain of the originator of a message or other object in the same way that a Notary Public assures a signature on a legal document. This is done by means of a Public Key cryptosystem. A public key cryptosystem differs from a private key cryptosystem in several ways. First, a public key system uses two keys, a Public Key that anyone can use (hence the name) and a private key that only the originator of a message uses. The originator uses the private key to encrypt the message (as in DES). The receiver, who has obtained the public key for the originator, may then decrypt the message. In this scheme, the public key is used to authenticate the originator′s use of his

or her private key, and hence the identity of the originator is more rigorously proven. The most widely known implementation of a public key cryptosystem is the RSA system. The Internet standard for privacy enhanced mail makes use of the RSA system.

8.3.5.8 Information Integrity

Information integrity refers to the state of information such that it is complete, correct, and unchanged from the last time in which it was verified to be in an integral state. The value of information integrity to a site will vary. For example, it is more important for military and government installations to prevent the disclosure of classified information, whether it is right or wrong. A bank, on the other hand, is far more concerned with whether the account information maintained for its customers is complete and accurate. Numerous computer system mechanisms, as well as procedural controls, have an influence on the integrity of system information. Traditional access control mechanisms maintain controls over who can access system information. These mechanisms alone are not sufficient in some cases to provide the degree of integrity required. Some other mechanisms are briefly discussed below. It should be noted that there are other aspects to maintaining system integrity besides these mechanisms, such as two-person controls, and integrity validation procedures.

Checksums: Easily the simplest mechanism, a simple checksum routine can compute a value for a system file and compare it with the last known value. If the two are equal, the file is probably unchanged. If not, the file has been changed by some unknown means. Though it is the easiest to implement, the checksum scheme suffers from a serious failing in that it is not very sophisticated and a determined attacker could easily add enough characters to the file to eventually obtain the correct value. A specific type of checksum, called a CRC checksum, is considerably more robust than a simple checksum. It is only slightly more difficult to implement and provides a better degree of catching errors. It too, however, suffers from the possibility of compromise by an attacker. Checksums may be used to detect the altering of information. However, they do not actively guard against changes being made. For this, other mechanisms such as access controls and encryption should be used.

Cryptographic Checksums: Cryptographic checksums (also called cryptosealing) involve breaking a file up into smaller chunks, calculating a (CRC) checksum for each chunk, and adding the CRCs together. Depending upon the exact algorithm used, this can result in a nearly unbreakable method of determining whether a file has been changed. This mechanism suffers from the fact that it is sometimes computationally intensive and may be prohibitive except in cases where the utmost integrity protection is desired. Another related mechanism, called a one-way hash function (or a manipulation detection code (MDC)) can also be used to uniquely identify a file. The idea behind these functions is that no two inputs can produce the same output, thus a modified file will not have the same hash value. One-way hash functions can be implemented efficiently on a wide variety of systems, making unbreakable integrity checks possible. (Snefru, a one-way hash function available via USENET as well as the Internet is just one example of an efficient one-way hash function.)

8.3.5.9 Limiting Network Access

The dominant network protocols in use on the Internet, carry certain control information that can be used to restrict access to certain hosts or networks within an organization. The IP packet header contains the network addresses of both the sender and recipient of the packet. Further, the TCP and UDP protocols provide the notion of a port, which identifies the endpoint (usually a network server) of a communications path. In some instances, it may be desirable to deny access to a specific TCP or UDP port, or even to certain hosts and networks altogether.

Gateway Routing Tables: One of the simplest approaches to preventing unwanted network connections is to simply remove certain networks from a gateway′s routing tables. This makes it impossible for a host to send packets to these networks. (Most protocols require bidirectional packet flow even for unidirectional data flow, thus breaking one side of the route is usually sufficient.) This approach is commonly taken in firewall systems by preventing the firewall from advertising local routes to the outside world. The approach is deficient in that it often prevents too much. In order to prevent access to one system on the network, access to all systems on the network is disabled.

Router Packet Filtering: Many commercially available gateway systems (more correctly called routers) provide the ability to filter packets based not only on sources or destinations, but also on source destination combinations. This mechanism can be used to deny access to a specific host, network, or subnet from any other host, network, or subnet. Gateway systems from some vendors support an even more complex scheme, allowing finer control over source and destination addresses. Via the use of address masks, one can deny access to all but one host on a particular network. Source routed packets may be filtered out by gateways, but this may restrict other legitimate activities, such as diagnosing routing problems.

8.3.5.10 Authentication Systems

Authentication refers to the process of proving a claimed identity to the satisfaction of some permission-granting authority. Authentication systems are hardware, software, or procedural mechanisms that enable a user to obtain access to computing resources. At the simplest level, the system administrator who adds new user accounts to the system is part of the system authentication mechanism. At the other end of the spectrum, fingerprint readers or retinal scanners provide a very high-tech solution to establishing a potential user′s identity. Without establishing and proving a user′s identity prior to establishing a session, your site′s computers are vulnerable to any sort of attack. Typically, a user authenticates himself or herself to the system by entering a password in response to a prompt. Challenge/response mechanisms improve upon passwords by prompting the user for some piece of information shared by both the computer and the user (such as mother′s maiden name, etc.).

Kerberos: Kerberos, named after the dog who in mythology is said to stand at the gates of Hades, is a collection of software used in a large network to establish a user′s claimed identity. Developed at the Massachusetts Institute of Technology (MIT), it uses a combination of encryption and distributed databases so that a user at a campus facility can log in and start a session from any computer located on the campus. This has clear advantages in certain environments where there are a large number of potential users who

may establish a connection from any one of a large number of workstations. Some vendors are now incorporating Kerberos into their systems.

Smart Cards: Several systems use smart cards (a small calculator-like device) to help authenticate users. These systems depend on the user having an object in their possession. One such system involves a new password procedure that requires a user to enter a value obtained from a smart card when asked for a password by the computer. Typically, the host machine will give the user some piece of information that is entered into the keyboard of the smart card. The smart card will display a response which must then be entered into the computer before the session will be established. Another such system involves a smart card which displays a number which changes over time, but which is synchronized with the authentication software on the computer.

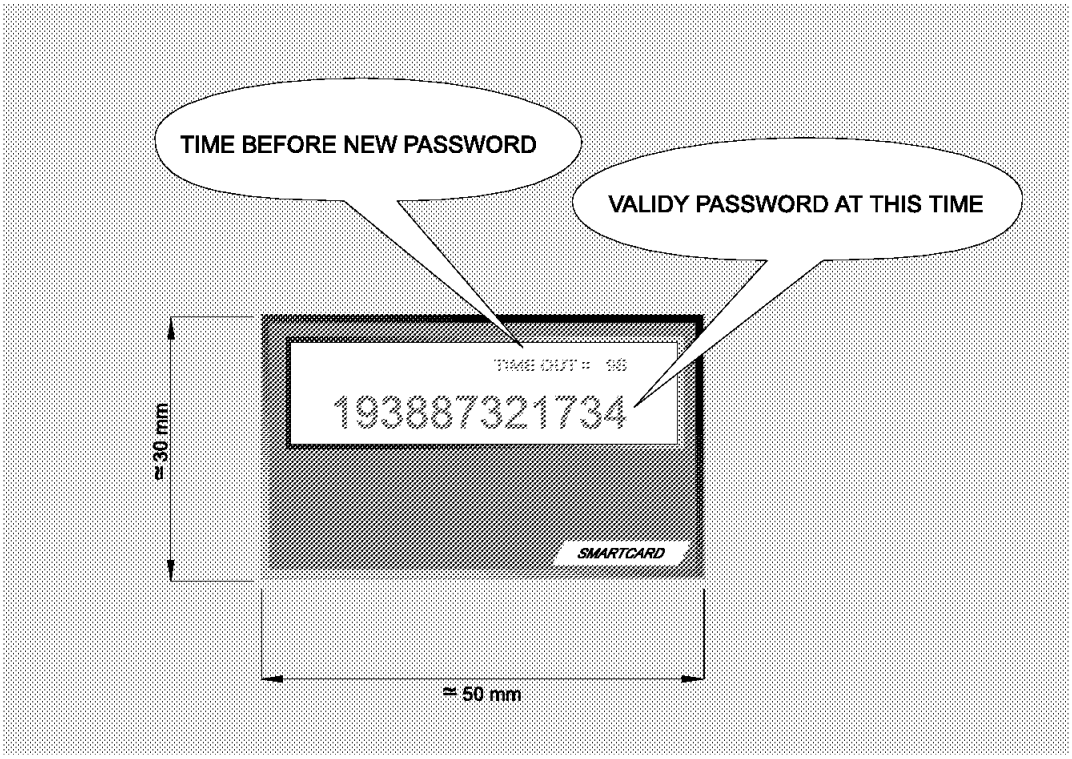


Figure 170. Smart card. The password synchronized smart card.

This is a better way of dealing with authentication than with the traditional password approach. On the other hand, some say it's inconvenient to carry the smart card. Startup costs are likely to be high as well.

Books, Lists, and Informational Sources: There are many good sources for information regarding computer security. The annotated bibliography at the end of this redbook can provide you with a good start. In addition, information can be obtained from a variety of other sources, some of which are described in this section.

8.3.6 Problem Reporting

8.3.6.1 Auditing

Auditing is an important tool that can be used to enhance the security of your installation. Not only does it give you a means of identifying who has accessed your system (and may have done something to it) but it also gives you an indication of how your system is being used (or abused) by authorized users and attackers alike. In addition, the audit trail traditionally kept by computer systems can become an invaluable piece of evidence should your system be penetrated.

Verify Security: An audit trail shows how the system is being used from day to day. Depending upon how your site audit log is configured, your log files should show a range of access attempts that can show what normal system usage should look like. Deviation from that normal usage could be the result of penetration from an outside source using an old or stale user account. Observing a deviation in logins, for example, could be your first indication that something unusual is happening.

Verify Software Configurations: One of the ruses used by attackers to gain access to a system is by the insertion of a so-called trojan horse program. A trojan horse program can be a program that does something useful, or merely something interesting. It always does something unexpected, like steal passwords or copy files without your knowledge. Imagine a trojan login program that prompts for a user name and password in the usual way, but also writes that information to a special file that the attacker can come back and read at will. Imagine a trojan editor program that, despite the file permissions you have given your files, makes copies of everything in your directory space without you knowing about it.

This points out the need for configuration management of the software that runs on a system, not as it is being developed, but as it is in actual operation. Techniques for doing this range from checking each command every time it is executed against some criterion (such as a cryptoseal, described above) or merely checking the date and time stamp of the executable. Another technique might be to check each command in batch mode at midnight.

8.3.7 Secure Web Servers

The World Wide Web (WWW) is a distributed hypermedia system which is rapidly gaining acceptance among Internet users. Although many WWW browsers support other, pre-existing Internet application protocols, the native and primary protocol used between WWW clients and servers is the HyperText Transfer Protocol. The ease of use of the Web has prompted widespread interest in its employment as a client/server architecture for many applications. Many such applications require the client and server to be able to authenticate each other and exchange sensitive information confidentially. Current HTTP implementations have only modest support for the cryptographic mechanisms appropriate for such transactions. Secure HTTP (S-HTTP) and Secure Socks Layer are special protocols that provides secure communication mechanisms between the browser and the server in order to enable spontaneous commercial transactions for a wide range of applications.

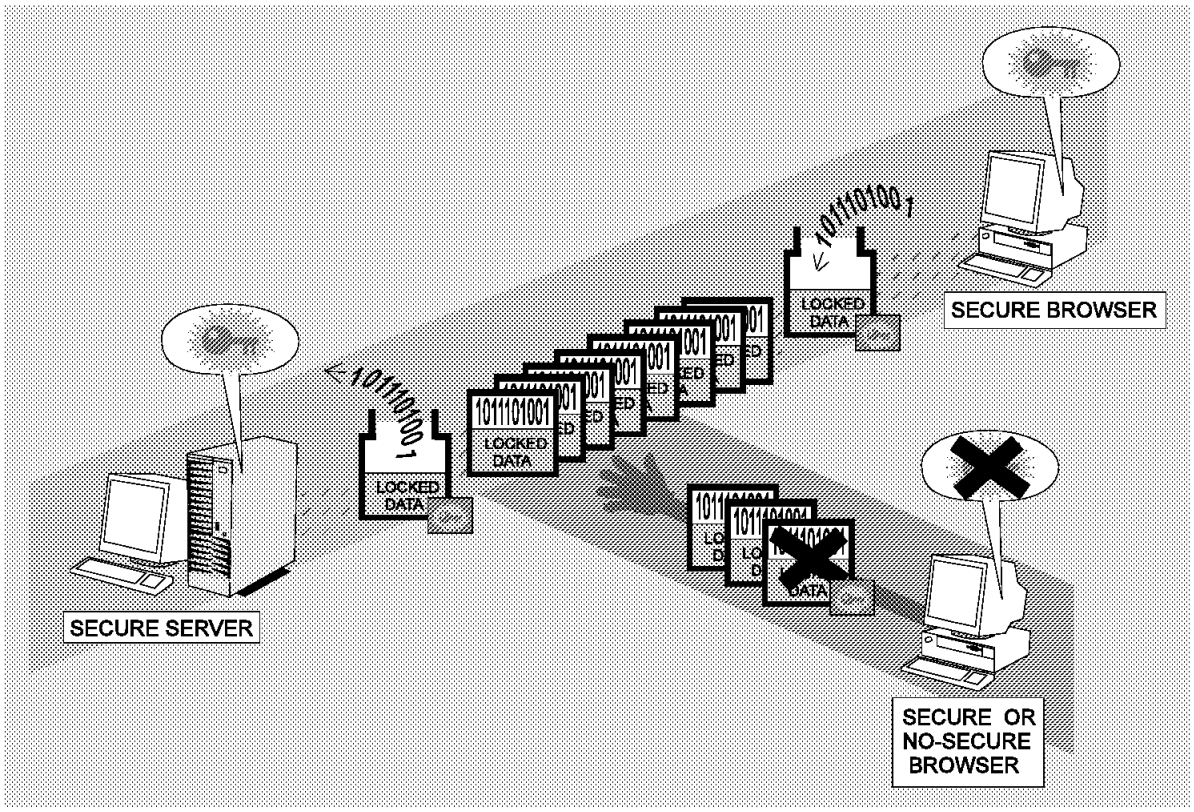


Figure 171. Secure Web Server. All data is encapsulated using a secure protocol and sent across the TCP/IP channel. Only the server and the relative client at this moment can understand the data built in this secure protocol.

8.3.7.1 Secure Hypertext Transfer Protocol / S-HTTP

Secure HTTP (S-HTTP) provides secure communication mechanisms between an HTTP client/server pair in order to enable spontaneous commercial transactions for a wide range of applications.

Our design intent is to provide a flexible protocol that supports multiple orthogonal operation modes, key management mechanisms, trust models, cryptographic algorithms and encapsulation formats through option negotiation between parties for each transaction.

Secure HTTP supports a variety of security mechanisms to HTTP clients and servers, providing the security service options appropriate to the wide range of potential end uses possible for the World Wide Web. The protocol provides symmetric capabilities to both client and server (in that equal treatment is given to both requests and replies, as well as for the preferences of both parties) while preserving the transaction model and implementation characteristics of the current HTTP. Several cryptographic message format standards may be incorporated into S-HTTP clients and servers, including, but not limited to, PKCS-7, PEM, and PGP.

S-HTTP supports interoperation among a variety of implementations, and is compatible with HTTP. S-HTTP aware clients can talk to S-HTTP oblivious servers and vice versa, although such transactions obviously would not use S-HTTP security features.

S-HTTP does not require client-side public key certificates (or public keys), supporting symmetric session key operation modes. This is significant because it means that spontaneous private transactions can occur without requiring individual users to have an established public key. While S-HTTP will be able to take advantage of ubiquitous certification infrastructures, its deployment does not require it.

S-HTTP supports end-to-end secure transactions, in contrast with the existing de-facto HTTP authorization mechanisms which require the client to attempt access and be denied before the security mechanism is employed. Clients may be primed to initiate a secure transaction (typically using information supplied in an HTML anchor); this may be used to support encryption of fill-out forms, for example.

With S-HTTP, no sensitive data need ever be sent over the network in the clear. S-HTTP provides full flexibility of cryptographic algorithms, modes and parameters. Option negotiation is used to allow clients and servers to agree on transaction modes. Should the request be signed? Encrypted? Both? What about the reply?

S-HTTP attempts to avoid presuming a particular trust model, although its designers admit to a conscious effort to facilitate multiply- rooted hierarchical trust, and anticipate that principals may have many public key certificates.

Message protection may be provided on three orthogonal axes: signature, authentication, and encryption. Any message may be signed, authenticated, encrypted, or any combination of these (including no protection).

8.3.7.2 Secure Socks Layer

The SSL protocol is designed to provide privacy between two communicating applications (a client and a server). Second, the protocol is designed to authenticate the server, and optionally the client. SSL requires a reliable transport protocol for data transmission and reception. The advantage of the SSL protocol is that it is application protocol independent. A higher level application protocol (for example. HTTP, FTP, TELNET, etc.) can layer on top of the SSL protocol transparently. The SSL protocol can negotiate an encryption algorithm and session key as well as authenticate a server before the application protocol transmits or receives its first byte of data. All of the application protocol data is transmitted encrypted, ensuring privacy. The SSL protocol provides channel security which has three basic properties:

- The channel is private. Encryption is used for all messages after a simple handshake is used to define a secret key.
- The channel is authenticated. The server endpoint of the conversation is always authenticated, while the client endpoint is optionally authenticated.
- The channel is reliable. The message transport includes a message integrity check (using a MAC).

In SSL, all data sent is encapsulated in a *record*, an object which is composed of a header and some non-zero amount of data. The primary goal of the SSL protocol is to provide privacy and reliability between two communicating applications. The protocol is composed of two layers. At the lowest level, layered on top of some reliable transport protocol is the SSL

Record Protocol. The SSL Record Protocol is used for encapsulation of various higher level protocols. One such encapsulated protocol, the SSL Handshake Protocol, allows the server and client to authenticate each other and to negotiate an encryption algorithm and cryptographic keys before the application protocol transmits or receives its first byte of data. One advantage of SSL is that it is application protocol independent. A higher level protocol can layer on top of the SSL Protocol transparently. The SSL protocol provides connection security that has three basic properties:

- The connection is private. Encryption is used after an initial handshake to define a secret key. Symmetric cryptography is used for data encryption.
- The peer's identity can be authenticated using asymmetric, or public key, cryptography.
- The connection is reliable. Message transport includes a message integrity check using a keyed MAC. Secure hash functions (for example, SHA, MD5, etc.) are used for MAC computations.

The goals of SSL Protocol, in order of their priority, are:

- **Cryptographic security:** SSL should be used to establish a secure connection between two parties.
- **Interoperability:** Independent programmers should be able to develop applications utilizing SSL that will then be able to successfully exchange cryptographic parameters without knowledge of one another's code.
- **Extensibility:** SSL seeks to provide a framework into which new public key and bulk encryption methods can be incorporated as necessary. This will also accomplish two sub-goals: to prevent the need to create a new protocol (and risking the introduction of possible new weaknesses) and to avoid the need to implement an entire new security library.
- **Relative efficiency:** Cryptographic operations tend to be highly CPU-intensive, particularly public key operations. For this reason, the SSL protocol has incorporated an optional session caching scheme to reduce the number of connections that need to be established from scratch. Additionally, care has been taken to reduce network activity.

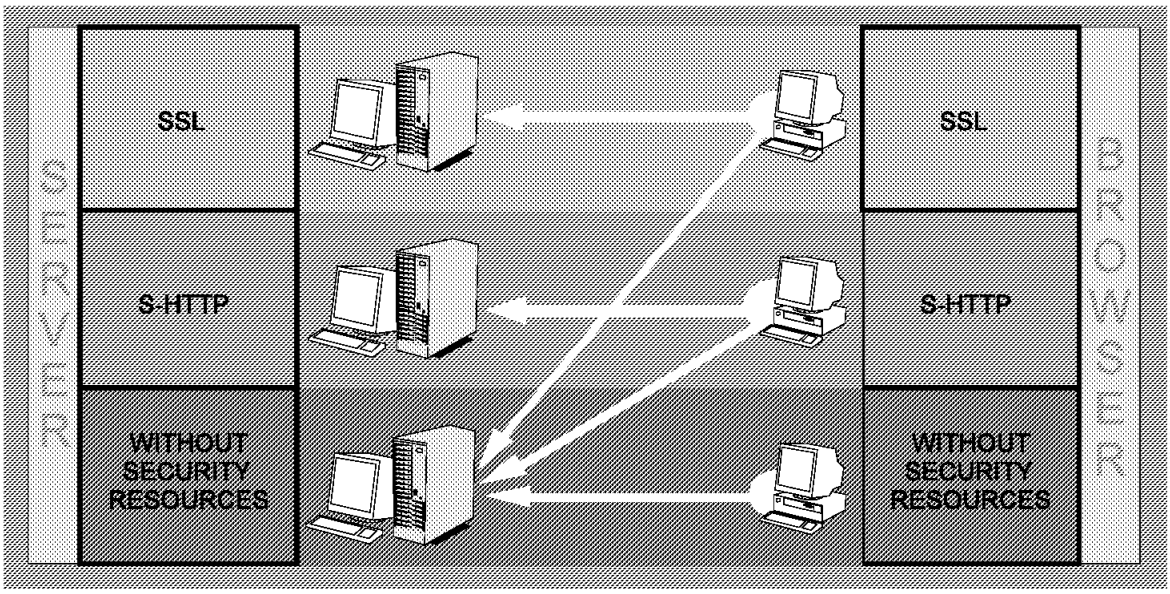


Figure 172. SSL and S-HTTP Protocols. The browsers that supports SSL and HTTP can access servers that are not using security resources, but the non-secure browsers cannot access this secure server when the security resources are enabled.

8.3.8 IBM Internet Connection Secure Products

The IBM Internet Connection Servers and Secure WebExplorer provide security resources using the S-HTTP and SSL technologies. Both protocols are supported on the servers and on the WebExplorer. The IBM Internet Connection Secure Servers and Secure WebExplorer browsers not only support SSL and S-HTTP, they also support a protocol called HTTPS that allows HTML documents to link to SSL-protected documents. HTTPS links can be specified in an anchor to protected documents or client users can code the reference directly by prefixing the document name with https://. Since HTTPS and HTTP are different protocols and use different ports, administrators can run secure and non-secure HTTP servers at the same time. This approach allows companies to offer catalog information to anyone while protecting themselves and clients during order entry. This offers the freedom, flexibility, and efficiency of HTTP while using SSL to protect sensitive parts of a transaction.

The IBM Internet Connection Secure Servers are available for OS/2, WindowsNT, AIX, MVS, Sun Solaris and HP-UX. Using these servers, you can:

- Distribute a wealth of up-to-date presale or postsale information to the world, using text, high-quality graphics, and even audio and video
- Create information that your customers and suppliers can interact with through electronic forms or e-mail
- Publish product descriptions and price lists with electronic order forms so your customers can purchase your product or service using a credit card, right from their computers
- Track how your customers, suppliers, and personnel use the information you publish so you can tell when you are reaching your target audience
- Provide all services listed above using security technologies

Other available features on the IBM Internet Connection Secure Servers are:

- Can be accessed by any industry-standard browser
- Can be easily installed, configured and used
- Tested extensively to ensure reliable operation
- Backed by IBM worldwide service and support
- Enabled for national language support

Table 29. IBM Internet Connection Secure Products

Service	IBM Product	Available operating system
Firewall	IBM Secure Network Gateway	• IBM AIX
Web Server	IBM Internet Connection Secure Servers	• OS/2 Warp • AIX • WindowsNT • Sun Solaris • HP-UX • MVS
Browser	IBM WebExplorer	• OS/2 Warp • AIX

8.3.9 Eletronic Commerce

Using the Internet to conduct business involving the exchange of money is called electronic commerce. Two consortia have proposed extensions to SSL and S-HTTP for electronic commerce. These extensions, currently in draft form, have been submitted for comments. One consortium, of which IBM is a member, has chosen to build commerce-specific extensions on top of already widespread protocols like SSL and S-HTTP. The other, led by Microsoft, has chosen to replace SSL and S-HTTP with its own protocols.

8.3.9.1 Electronic Money (e-money)

Public-key cryptography and digital signatures (both blind and non-blind signatures) make e-money possible. It would take too long to go into detail how public-key cryptography and digital signatures work. But the basic idea is that banks and customers would have public-key encryption keys. Public-key encryption keys come in pairs: a private key known only to the owner, and a public key, made available to everyone. Whatever the private key encrypts, the public key can decrypt, and vice versa. Banks and customers use their keys to encrypt (for security) and sign (for identification) blocks of digital data that represent money orders. A bank signs money orders using its private key and customers and merchants verify the signed money orders using the bank′s widely published public key. Customers sign deposits and withdrawals using their private key and the bank uses the customer′s public key to verify the signed withdrawals and deposits.

The different kinds of e-money: In general, there are two distinct types of e-money:

- Identified e-money and anonymous e-money (also known as digital cash). Identified e-money contains information revealing the identity of the

person who originally withdrew the money from the bank. Also, in much the same manner as credit cards, identified e-money enables the bank to track the money as it moves through the economy.

- Anonymous e-money works just like cash. Once anonymous e-money is withdrawn from an account, it can be spent or given away without leaving a transaction trail. You create anonymous e-money by using blind signatures rather than non-blind signatures.

There are two varieties of each type of e-money:

- Online e-money
- Offline e-money

Online means you need to interact with a bank (via modem or network) to conduct a transaction with a third party. Offline means you can conduct a transaction without having to directly involve a bank. Offline anonymous e-money (true digital cash) is the most complex form of e-money because of the double-spending problem.

The double-spending problem: Since e-money is a bunch of bits, a piece of e-money is very easy to duplicate. Since the copy is indistinguishable from the original you might think that counterfeiting would be impossible to detect. A trivial e-money system would allow us to copy of a piece of e-money and spend both copies. We could become millionaires in a matter of a few minutes. Obviously, real e-money systems must be able to prevent or detect double spending.

Online e-money systems prevent double spending by requiring merchants to contact the bank′s computer with every sale. The bank computer maintains a database of all the spent pieces of e-money and can easily indicate to the merchant if a given piece of e-money is still spendable. If the bank computer says the e-money has already been spent, the merchant refuses the sale. This is very similar to the way merchants currently verify credit cards at the point of sale.

Offline e-money systems detect double spending in a couple of different ways. One way is to create a special smart card containing a tamper-proof chip called an observer (in some systems). The observer chip keeps a mini database of all the pieces of e-money spent by that smart card. If the owner of the smart card attempts to copy some e-money and spend it twice, the imbedded observer chip would detect the attempt and would not allow the transaction. Since the observer chip is tamper-proof, the owner cannot erase the mini-database without permanently damaging the smart card.

The other way offline e-money systems handle double spending is to structure the e-money and cryptographic protocols to reveal the identity of the double spender by the time the piece of e-money makes it back to the bank. If users of the offline e-money know they will get caught, the incidence of double spending will be minimized (in theory). The advantage of these kinds of offline systems is that they don′t require special tamper-proof chips. The entire system can be written in software and can run on ordinary PCs or cheap smart cards.

It is easy to construct this kind of offline system for identified e-money. Identified offline e-money systems can accumulate the complete path the

e-money made through the economy. The identified e-money grows each time it is spent. The particulars of each transaction are appended to the piece of e-money and travel with it as it moves from person to person, merchant to vender. When the e-money is finally deposited, the bank checks its database to see if the piece of e-money was double spent. If the e-money was copied and spent more than once, it will eventually appear twice in the spent database. The bank uses the transaction trails to identify the double spender.

Offline anonymous e-money (sans observer chip) also grows with each transaction, but the information that is accumulated is of a different nature. The result is the same however. When the anonymous e-money reaches the bank, the bank will be able to examine its database and determine if the e-money was double spent. The information accumulated along the way will identify the double spender.

The big difference between offline anonymous e-money and offline identified e-money is that the information accumulated with anonymous e-money will only reveal the transaction trail if the e-money is double spent. If the anonymous e-money is not double spent, the bank can not determine the identity of the original spender nor can it reconstruct the path the e-money took through the economy.

With identified e-money, both offline or online, the bank can always reconstruct the path the e-money took through the economy. The bank will know what everyone bought, where they bought it, when they bought it, and how much they paid. And what the bank knows, the taxation authority knows.

There are a lot of companies developing products based on the e-money technology. They are:

- Cybercash/ www
- CheckFree
- Digicash
- First Virtual
- Netbill Project
- Software Agent′s Netbank
- USC′s Netcash
- NetCheque
- NetMarket
- Mondex
- GTE/ www.gte.com
- Master Card/ www.mastercard.com
- Netscape/ www.netscape.com
- Security First Network Bank, FSB/
- Visa/ www.visa.com
- IBM Corporation/ www.ibm.com
- Sandia′s Eletronic Cash System

- First Union Bank/ www.firstunion.com

8.3.9.2 Secure Electronic Payment Protocol

IBM, Netscape, GTE, CyberCash, and Master Card have cooperatively developed extensions they call the Secure Electronic Payment Protocol (SEPP). IBM has contributed both security technology including iKP (a secure payment technology developed at IBM′s research laboratory in Zurich, Switzerland) and its long-standing experience building and operating very large financial networks. SEPP protects transactions between a card holder and a merchant, and between the merchant and card holder′s financial institution. There are seven major business requirements addressed by the Secure Electronic Payment Protocol (SEPP) system:

- Confidentiality of payment information.
- Integrity of all payment data transmitted via public networks.
- Authentication that a card holder is the legitimate owner of a credit card account.
- Authentication that a merchant can accept credit card payments with an acquiring member financial institution.
- Interoperability of bank card/credit card programs among software and network providers.
- Protection from electronic commerce-related attacks.
- Separate privacy mechanisms for general information exchange and payment data exchange.

The SEPP system automates the highly manual system used today. In the SEPP system, the card holder begins the transaction sequence by sending the merchant a message. The merchant responds with a message containing transaction information used by the card holder. The card holder then prepares a request with encrypted order validation information and the card holder′s payment instructions. The merchant receives the request and passes it to the financial institution for confirmation. The financial institution processes the request and responds to the merchant with an authorization. The merchant responds to the card holder.

The process of shopping is set individually by merchants providing the service.

The process of transaction capture, clearing and settlement of the transaction, is defined by the relationship between the merchant and their financial institution.

The scope of SEPP encompasses both interactive on-line and non-interactive store-and-forward (e-mail message based) payment transactions. Several transaction messages are required; others add the ability to operate when the customer or the financial institution are not available. Card holder account and payment data information must be secured as it travels across the network, preventing interception and alteration of this data by unauthorized parties. The SEPP standard guarantees that message content is not altered during transmission. Payment data sent from card holders to merchants is protected in such a manner as to be verifiable. If any component is altered in transit, the transaction will not be processed accurately. SEPP provides the means to ensure that the contents of all

payment messages sent match the contents of messages received. Merchants will be able to verify that a card holder is using a valid account number.

A mechanism that links a card holder to a specific account number reduces the incidence of fraud and therefore the overall cost of payment processing.

SEPP also provides a mechanism to prevent intruders from establishing a phony storefront and collecting payment data. Merchants who receive payment data are sponsored by a financial institution and display a certificate verifying this relationship.

8.3.9.3 IBM Corporation iKP (Internet Keyed Payment Protocols)

The IBM Research Division has developed a family of secure payment protocols, called iKP that circumvent most of the above problems. While developed at IBM, the technology has been immediately disclosed for public review, and it is being openly discussed in a number of fora and consortia (for example, W3C,FSTC,IETF, etc.) and with a number of financial and technical partners as IBM has no intention of keeping it proprietary. The technology uses strong cryptography in a very secure way but packages it so that it should satisfy usage and import/export restrictions in most countries. It was designed to work with any browser and server on any platform; the first prototype of it is designed to work with credit cards, but the intrinsic design is flexible and will allow supporting other payment instruments in due time. This first prototype is also entirely in software because typical Internet stations today do not include secure hardware or support smart card readers, but provisions are made in the design to accommodate such devices later, and work is already in progress in that direction. The iKP technology is designed to allow customers to order goods, services, or information over the Internet, while relying on existing secure financial networks to implement the necessary payments, as suggested in the next figure.

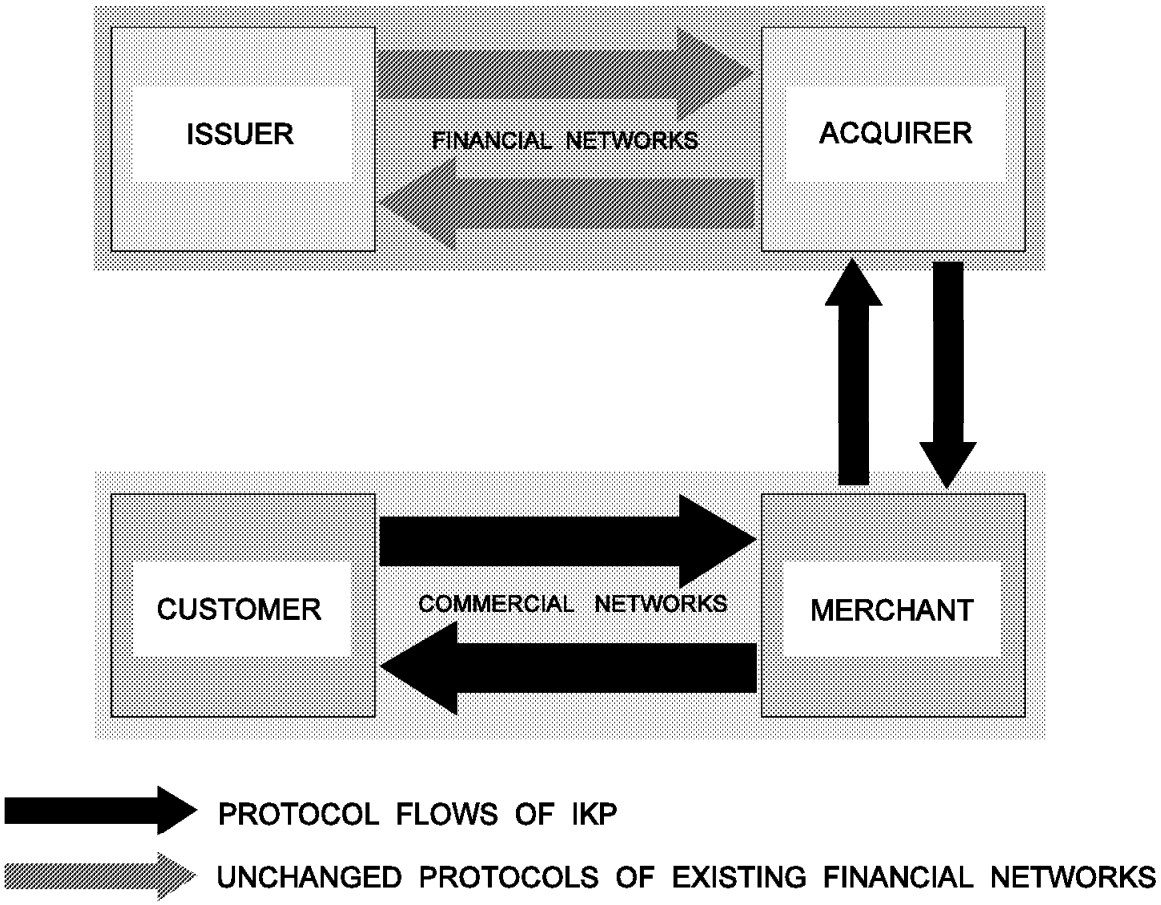


Figure 173. IBM iKP

The iKP technology is based on RSA public-key cryptography. Depending on requirements, an electronic payment transaction using iKP may involve one, two, or three public keys; in all cases the bank acquiring the transaction for processing will have a public-private key pair for receiving confidential information such as credit card numbers and signing authorization messages. In many cases the merchant will also have a public-private key pair for receiving confidential information and signing payment requests and purchase confirmations. In some cases even customers may have a public-private key pair for signing payment transactions. In all cases they have a PIN for confirming authorization of payment.

Certificate Management: The iKP technology is based on public-key cryptography (for example, RSA.) Depending on requirements, an electronic payment transaction using iKP may involve one, two, or three public keys. In all cases, the acquirer has a public-private key pair for receiving confidential information such as buyer account numbers and for signing authorization messages. Sellers may also have key pairs for signing payment requests and purchase confirmations. Buyers can have key pairs for signing (authorizing) payment transactions. The acquirer is the only entity that both signs and receives confidential data. An acquirer may have two public/private key pairs: one for signatures and one for encryption. However, both key pairs may be validated by a single certificate. The recipient of any signed message must hold a copy of the public key required

to validate the signature. Specifically, seller and buyer must both have a copy of the acquirer′s public key in order to validate the acquirer′s signature of the authorization method. The buyer also needs a copy of a different public key of the acquirer for encrypting the account number and related information. If the seller′s signature of the invoice is implemented, both buyer and acquirer need to have the seller′s public key. If the buyer signature of the payment is implemented, the acquirer (and, sometimes, the seller) needs to have the buyer′s public key.

Public keys are distributed to the participants in the form of certificates signed by some authority. Certificates can be distributed in two ways:

1. Before executing iKP, for example, during browsing or out-of-band, or
2. In the course of iKP execution, as part of iKP option fields.

In the former case, certificates may be cached from previous payment transactions, provided as part of HTML fields, transmitted via electronic mail, or communicated by any other means desired. Such mechanisms are outside the definition of iKP. The establishment of the certificate authority, and the communication of the authority′s root public key is also outside this protocol.

One possible design is for each credit card system to have a certificate authority with a well-known root public key. This authority would sign certificates for all acquirers, sellers, and buyers who utilize the credit card system. Alternatively, some other well-trusted organization could issue certificates for any or all iKP participants.

Any purchase transaction involves (at least) three phases:

1. Negotiation of the purchase terms and other details
2. Actual payment
3. Order fulfillment/delivery

The iKP is the electronic equivalent of the paper charge slip, signature, and submission process, or of a paper check with online funds verification. It comes after the negotiation is completed. iKP takes input from the negotiation process (payment amount, order description, payment method, etc.) and causes the payment to happen via a three-way communication among the buyer, seller, and acquirer. Negotiation is a bilateral conversation between the buyer and seller that may be implemented in many ways, for example, via HTTP using a WWW browser and server, electronic mail, paper catalog for the offer from the seller and electronic mail for the order from the buyer. The negotiation process addresses not only what is ordered (x units of these widgets and y units of those) but the terms of the order (prices, delivery addresses, schedules, credit card type), and the method of payment (cash, paper check, digital cash, iKP, whether a receipt is required, etc.). Irrespective of the means used to conduct negotiation, the buyer, at some point, initiates payment. This is the point when negotiation ends and iKP starts. The data required by iKP in the buyer system are: acquirer′s public key, seller′s public key (if implemented), buyer′s account number (BAN in the protocol description, see below), buyer′s public/private key pair (if implemented), buyer′s PIN (if implemented), payment amount and currency ($$), and the description of the order (DESC).

The data required by iKP in the seller system is: acquirer′s public key, seller′s ID, seller′s public/private key (if implemented), payment amount and currency ($$), and the description of the order (DESC).

From the perspective of iKP, order description (DESC) is an opaque string that is incorporated via a hash into the protocol to bind the description to the payment. Opaque means that iKP does not interpret the contents of the description. The only requirement of iKP is for the description to contain all relevant details of the transaction (ordered goods, delivery address, payment terms, etc.), and that both buyer and seller possess exactly the same opaque string.

iKP As an Architecture: iKP is a general architecture that accommodates a variety of payment method interactions by making certain message flows and fields optional. This document defines what types of security are supported by various combinations of options. Any particular use of iKP (for example, iKP for credit cards) will require a detailed specification for that particular use. iKP is intended for use with a number of different communications channels among the participants, for example, HTTP, SHTTP, and (electronic mail. Applications of the iKP architecture to specific communications environments are not discussed in this document. It is envisaged that other documents will define the syntax of iKP for each desired communications method. Hopefully there will be one syntax for each communications channel regardless of the purchase style (for example credit card versus debit card).

Fault Tolerance and Exception Handling: As can be expected in any communication environment, especially, in the Internet, absolute reliability is next to impossible. Therefore, in order to design, not only secure, but also robust, payment protocols, we need to consider all possible anomalous scenarios. No assumptions are made below about the robustness of the underlying network infrastructure since it is envisaged that the iKP protocol will operate in environments with widely varying degrees of reliability. It is assumed that all parties in iKP (except acquirer) implement timeouts and retransmissions whenever a message elicits no reply. All unexpected messages, for example, those not corresponding to an outstanding or recorded transaction, are ignored. All invalid messages (for example, acquirer receiving INITIATE) are similarly ignored. The term duplicate is used to mean that the message is otherwise valid. Also, the term unsolicited is used to mean that the message is otherwise valid, for example, all contained signatures (if any) are verifiable. All parties are assumed to hae access to stable, non-volatile storage. The term recording is used to mean commitment to stable storage.

Refunds: Credit card systems support the concept of returns or refunds. The buyer returns merchandise to the seller along with the original credit card slip. The seller issues a refund slip which causes all or part of the original payment amount to be credited to the buyer′s credit card account. An analogous function can be achieved in iKP but only if the seller can sign.

To process a refund, buyer and seller simply run iKP using a negative amount, effectively crediting rather than debiting money to the buyer′s account. This may be repeated multiple times if the buyer returns portions of an order in multiple refund transactions. As an option, the seller and acquirer may require that the CONFIRM message from a purchase be

associated (in the optional TEXT fields) with any refund. This permits the seller and acquirer to validate the refund amount against the original purchase amount. It permits the seller to verify the original purchase transaction and detect multiple refunds that total to more than the original purchase.

Order Status Inquiry: Given the distinction between authorization and clearance, buyers may want a method of finding out from sellers whether a payment has cleared. This is one instance of many kinds of order status inquiry. For example, buyers may wish to know whether purchased goods have actually been shipped by the seller. Such inquiry functions are outside the scope of iKP because they are not required for payment, they involve bilateral (rather than multi-party) communication and they extend to a variety of non-payment issues.

Security Considerations: The intent of iKP is to address certain security issues related to three-party payment mechanisms conducted over the Internet. Note that iKP does not address security concerns applicable to negotiations that may occur before iKP is initiated. Depending upon the communications method utilized, security protocols such as SSL (2), SHTTP (3), PEM (4), or MOSS (5) should be utilized if privacy, authentication, signatures, or other security attributes are required for the negotiations.

Public key signature mechanisms are critically dependent upon the security of the corresponding private keys. iKP requires private and public keys of acquirers and optionally of sellers and buyers. Implementors should pay particular attention to the methods used to store the private keys of these participants. Encryption of stored private keys, tamper-proof hardware, certificate revocation mechanisms, and certificate expiration dates should all be considered. iKP expects that public keys are distributed via certificates signed by well-known certification authorities (CAs).

The definition of such CAs, and the distribution mechanism for their root public keys, is outside the scope of iKP. The security of iKP ultimately relies upon the security of the root keys as utilized by the buyer, seller, and acquirer software. Implementors should consider carefully how software configures and stores these root keys. It is suggested that there be mechanisms by which buyers, sellers, and acquirer employees/users can verify the certificate authorities and root keys recognized by their software.

8.3.9.4 Security Mailing Lists

The UNIX Security Mailing List exists to notify system administrators of security problems before they become common knowledge, and to provide security enhancement information. It is a restricted-access list, open only to people who can be verified as being principal systems people at a site. Requests to join the list must be sent by either the site contact listed in the Defense Data Network′s Network Information Center′s (DDN NIC) WHOIS database, or from the root account on one of the major site machines. You must include the destination address you want on the list, an indication of whether you want to be on the mail reflector list or receive weekly digests, the electronic mail address and voice telephone number of the site contact if it isn′t you, and the name, address, and telephone number of your organization. This information should be sent to SECURITY-REQUEST@CPD.COM.

The RISKS digest is a component of the ACM Committee on Computers and Public Policy. It is a discussion forum on risks to the public in computers and related systems, and along with discussing computer security and privacy issues, has discussed such subjects as the Stark incident, the shooting down of the Iranian airliner in the Persian Gulf (as it relates to the computerized weapons systems), problems in air and railroad traffic control systems, software engineering, and so on. To join the mailing list, send a message to RISKS-REQUEST@CSL.SRI.COM. This list is also available in the USENET newsgroup comp.risks.

The VIRUS-L list is a forum for the discussion of computer virus experiences, protection software, and related topics. The list is open to the public, and is implemented as a moderated digest. Most of the information is related to personal computers, although some of it may be applicable to larger systems. To subscribe, send the line:

```
SUB VIRUS-L your full name
```

to the address LISTSERV%LEHIIBM1.BITNET@MITVMA.MIT.EDU. This list is also available via the USENET newsgroup comp.virus.

8.3.9.5 Networking Mailing Lists

The TCP/IP Mailing List is intended to act as a discussion forum for developers and maintainers of implementations of the TCP/IP protocol suite. It also discusses network-related security problems when they involve programs providing network services, such as Sendmail. To join the TCP/IP list, send a message to TCP/IP-REQUEST@NISC.SRI.COM. This list is also available in the USENET newsgroup comp.protocols.tcp/ip. The USENET groups misc.security and alt.security also discuss security issues. misc.security is a moderated group and also includes discussions of physical security and locks. alt.security is unmoderated.

8.3.10 Reference Sites on the Internet

S-HTTP memo

http://www.commerce.net/information/standards/drafts/shttp.txt

Site Security Handbook

http://www.net.ohio-state.edu/hypertext/rfc1244/toc.html

SSL, S-HTTP and Security related links

http://www.netscape.com/newsref/std/index.html

Firewalls Reference

http://www.net.ohio-state.edu/faq/usenet/firewalls-faq/faq.html

General security documents

http://www.yahoo.com/Business_and_Economy/Companies/Computers/Security
http://www.sei.cmu.edu/SEI/programs/cert.html
http://mls.saic.com/mls.security.html http://everest.cs.ucdavis.edu
http://www.cs.purdue.edu/coast/coast.html

Chapter 9. Network Management

In this chapter we introduce network management as part of the SystemView architecture. We also introduce network management in the Internet environment and the *de facto* method of managing these networks SNMP. We finish the chapter with the products IBM offers in this area.

9.1 SystemView Introduction

In 1990, IBM announced the SystemView strategy for planning, coordinating and operating heterogeneous, enterprise-wide information systems. This strategy comprises the IBM SystemView structure and SystemView conforming products. SystemView is the SAA (Systems Application Architecture) strategy for managing enterprise information systems.

The SystemView structure is designed to provide system users with a consistent interface, shared data, enhanced automation and increased interaction among system management products. Products conforming to the SystemView structure provide management functions that span information systems resources in SAA environments as well as other IBM and non-IBM environments. These resources may be managed across OSI (Open Systems Interconnection), TCP/IP and SNA networks. This systems management strategy enhances the ability of users to manage enterprise-wide information systems as a business and to provide quality service to help achieve the goals of the enterprise.

SystemView addresses the management of the following resources:

- Hosts
- Databases
- Auxiliary storage
- Networks
- Business administration (of information systems)

SystemView provides end-to-end management solutions for both distributed and host systems environments. The flexibility provided by multiple managing systems - Operating System/2 (OS/2), Advanced Interactive Executive (AIX/6000), Operating System/400 (OS/400) and NetView from IBM makes it possible to extend system and network monitoring and control to AIX/6000-based, DOS-based and OS/2-based local area networks (LANs), as well as the Application System/400 (AS/400) family. This same capability can also be extended to products managing distributed and centralized data, text, voice, graphics and image information.

9.1.1 SystemView Benefits

The IBM SystemView management strategy provides:

- The SystemView structure for integrating systems management applications from IBM, outside vendors and IBM customers
- User productivity gains through the use of consistent user interfaces, standardized systems management data definitions, increased integration, and enhanced automation

© Copyright IBM Corp. 1996

- Enhanced business solutions as a result of increased flexibility and extendibility through the use of open standards
- Customer investment protection through an evolutionary approach and orderly migration paths
- Customer growth through the increased availability of systems and networks
- Increased level of automation for systems management tasks
- Architected interfaces to enable vendor and customer participation

9.1.2 SystemView Structure

The IBM systems management strategy consists of the SystemView structure and SystemView conforming products. SystemView structure consists of three complementary elements called *dimensions* which define guidelines, standards and interfaces for integrating systems management applications.

Table 30. SystemView Structure and Application Dimension Disciplines

S y s t e m V i e w Structure		
	Application Dimension	
	Business Management	
	Change Management	
End-Use Dimension	Configuration Management	**Data Dimension**
	Operations Management	
	Performance Management	
	Problem Management	

- The End-Use Dimension provides the user at a workstation with a consistent, user-friendly view of the applications.
- The Application Dimension defines guidelines for the implementation and integration of systems management applications.
- The Data Dimension addresses requirements for standardized systems management data definitions and access.

The End-Use Dimension: The End-Use Dimension addresses the needs of SystemView end users, such as the operators, system administrators, and business analysts who perform systems management tasks. The End-Use Dimension provides definitions for the presentation of systems management objects and actions. These definitions are designed to provide common semantics, appearance, behavior and terminology across related SystemView applications, thereby increasing end user productivity and reducing the overall required training effort.

The End-Use Dimension allows the user a choice of interfaces, such as graphic display, textual dialogs, or a command entry. Methods and interfaces are defined for use within SystemView applications, along with tools and services.

The Data Dimension: The Data Dimension provides the platform for integrating all systems management data in accordance with a data model defined by SystemView. Within this platform there are interfaces and services which can be used by applications seeking access to the systems management data.

The Data Dimension provides a common data model for systems management data. This prevents data redundancy and ensures consistency among the different systems and products.

The Application Dimension: The Application Dimension provides a comprehensive approach to integrating systems management tasks and applications. The Application Dimension defines the interfaces and services necessary to support the tasks required to administer, coordinate, and operate the enterprise systems as a business. These systems management tasks are called disciplines and are grouped into the following six management areas:

- Business management
- Change management
- Configuration management
- Operations management
- Performance management
- Problem management

Business management includes tasks that support a wide range of business and administrative functions to run the business aspects of enterprise-wide information systems. Examples of business management tasks are security management, inventory/asset control, accounting, billing and charge-back and budget planning.

Change management includes tasks that manage and control the introduction of change into a systems environment. These would include planning, testing and distribution of changes to data processing resources.

Configuration management is the collection of the facilities and processes needed to plan, develop and maintain the operational properties and interrelationships of resources within the enterprise′s information systems. The design and updating of configuration information are two of the tasks which fall into this category.

Operations management deals with tasks that plan, distribute, evaluate and control workloads. Examples are tasks which include workload and operations planning, scheduling and control.

Performance management addresses the effectiveness with which information systems deliver services to their customers. Service planning and control are examples of performance management tasks.

Problem management is the process of managing problems, incidents, and critical situations from their detection until their final resolution. Incident detection and recognition as well as problem analysis and diagnosis would be grouped under this discipline.

9.2 Managing a Heterogeneous Network

Today there are many manufacturers producing hundreds of devices such as personal computers, routers and mainframes which support TCP/IP. Due to the open nature of TCP/IP and the Internet, many networks have become heterogeneous and multivendor in makeup. Vendor-specific network management tools were found to be unusable in these environments. It became obvious that an open network management technology was required to manage these networks. Thus, SNMP has become the industry standard network management protocol for heterogeneous networks.

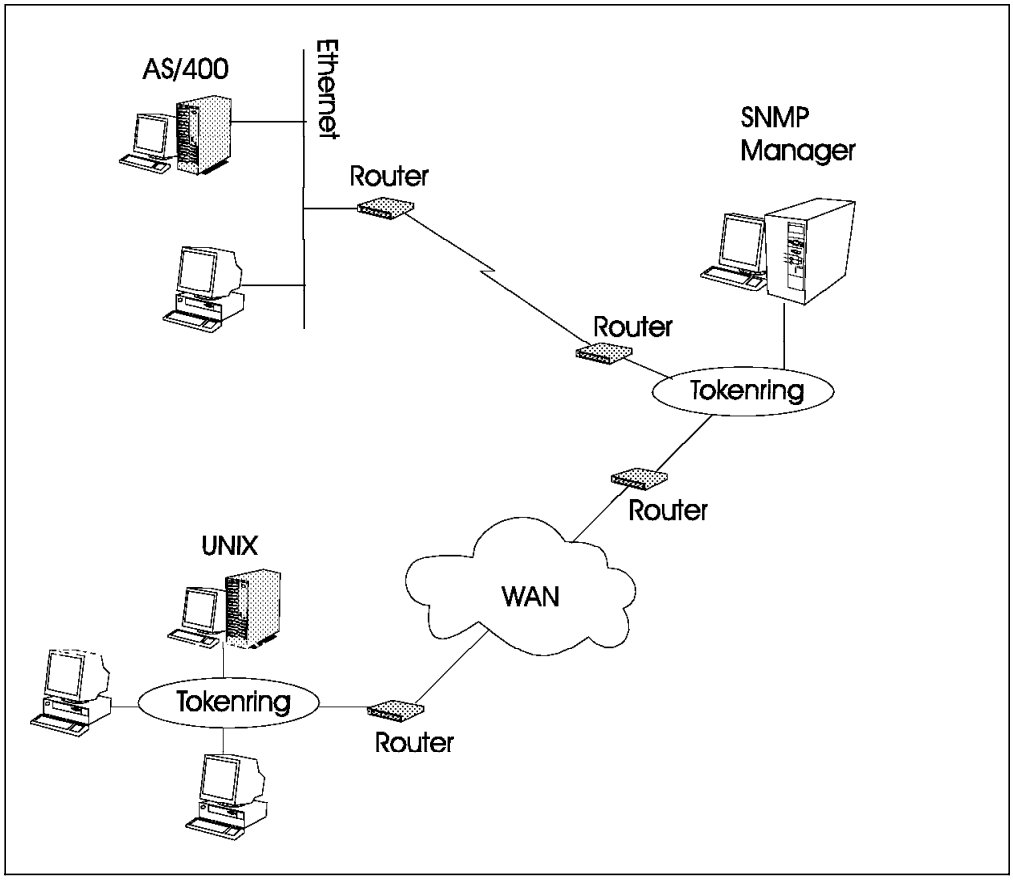


Figure 174. A Typical Heterogeneous Network

9.2.1 A Brief View into SNMP History

In 1968, the U.S. Defense Advanced Research Projects Agency (DARPA) began an effort to develop a technology which is now known as packet switching. This technology was strongly influenced by the development of low-cost minicomputers and digital telecommunications techniques during the 1960s. In the early 1970s, the DARPA sponsored several programs to explore the use of packet switching methods in alternative media such as mobile radio and satellite.

The expansion of the Internet drew support from U.S. government organizations including DARPA, the National Science Foundation (NSF), the Department of Energy (DOE), and the National Aeronautics and Space Administration (NASA). Eventually, international research bodies also got involved in the Internet.

Due to the successful implementation of packet radio and packet satellite technology, the desire to connect the DARPA network, ARPANET, with other packet nets arose. This led to the development of an internetwork protocol and a set of gateways to connect the different networks. DARPA sponsored further development of this solution, which resulted in a collection of computer communications protocols based on the original Transmission Control Protocol (TCP) and the lower level Internet Protocol (IP). During the course of the research, many other protocols were developed. These protocols, together with TCP and IP, are referred to as the TCP/IP Protocol Suite. A protocol suite is a set of protocols that work cooperatively together.

During these early stages, network management was of a proprietary nature due to the fact that networks were constructed with vendor-specific technology. In recognition of the need for a network management framework suitable for non-proprietary technology, in the late 1970s, the International Organization for Standardization (ISO), together with the International Telephone and Telegraph Consultative Committee (CCITT), started a research effort on this subject, resulting in the Open Systems Interconnection (OSI) protocol suite.

As the number of interconnected networks began to increase during the 1980s, the management of the Internet grew more complicated because the networks were using equipment from different vendors. In order to meet the network management demands at hand, the Internet Activities Board (IAB) defined a strategy formed by two parts:

- In the short term, the Simple Gateway Monitoring Protocol (SGMP) being of simpler nature than the OSI model, would be modified in order to produce a new protocol for managing nodes in the Internet community.
- In the long term, the network management protocol called Common Management Information Protocol (CMIP), used in the OSI model would continue to be observed.

The enhancements made to SGMP eventually originated SNMP. Currently, the simple network management protocol (SNMP) is an industry standard protocol which is used for network and system management. SNMP is a collection of specifications which describe how to manage and control a Network Element (SNMP agent) from a network managing station (SNMP manager). The SNMP specifications are contained in documents called Request for Comments (RFC), which are controlled by the IAB.

The RFCs that define the SNMP specifications are the following:

- RFC1155: Structure and identification of management information for TCP/IP-based Internets
- RFC1212: Concise MIB definition
- RFC1213: Management information base for network management of TCP/IP-based iternets: MIB-II
- RFC1157: Simple network management protocol (SNMP)

For further details about the IAB, and RFCs, see Appendix A, “The IAB” on page 559.

Although SNMP is used predominantly in TCP/IP-based networks, AnyNet sockets over SNA allows SNMP support to be used in SNA networks.

9.2.2 SNMP Definitions

SNMP is a transaction-oriented protocol that allows network elements to be queried directly. It is a simple protocol that allows management information for a network element to be inspected or altered by a system administrator at a network management station. SNMP is a TCP/IP network management protocol and is based on a manager and agent interaction. The SNMP manager (such as NetView for OS/2) communicates with its agents. Agents gather management data while the managers solicit this data and process it. An agent can also send unsolicited information, called a trap, to a managing system to inform it of an event that has taken place at the agent system. For example, an agent can send a trap of type linkDown to the manager to inform it about the loss of a communication link with a particular device.

SNMP Agent: An SNMP agent is an implementation of a network management application which is resident on a managed system. Each node that is to be monitored or managed by an SNMP manager in a TCP/IP network, must have an SNMP agent resident. The agent receives requests to either retrieve or modify management information by referencing MIB objects. MIB objects are referenced by the agent whenever a valid request from an SNMP manager is received.

SNMP Manager: An SNMP manager refers to a managing system that executes a managing application or suite of applications. These applications depend on MIB objects for information that resides on the managed systems.

SNMP Subagent: An SNMP subagent is the implementation of a network management application on a managed system, which interfaces with the SNMP agent for the purpose of expanding the number of MIB objects that an SNMP manager can access. SNMP agents have predefined MIB objects that they can access. This limits the managing application in regards to the type of information that it can request. The need to overcome this limitation brought about the introduction of subagents. A subagent allows the dynamic addition of other MIB objects without the need to change the agent. Whether a MIB object is referenced by the agent or the subagent is transparent to the managing system.

SNMP Proxy Agent: An SNMP proxy agent is one that acts on behalf of a managed system that is not reached directly by the managing system. A proxy agent is used when a managed system does not support SNMP, or when a managed system supports SNMP but for other reasons it is more convenient to manage it indirectly, for instance, through the use of a proxy agent.

Management Information Base (MIB): A management information base (MIB) is a logical database residing in the managed system which defines a set of MIB objects. A MIB is considered a logical database because actual data is not stored in it, but rather provides a view of the data that can be accessed on a managed system.

MIB Object A MIB object is a unit of managed information that specifically describes an aspect of a system, for example, CPU utilization, software name, hardware type, and more. A collection of related MIB objects is defined as a management information base (MIB). A MIB object is sometimes called a MIB variable.

Instance An instance refers to a particular representation of a MIB object. The MIB object which it represents can be thought of as a template for which one or more instances can be defined, depending on the type of MIB object. Actual values can only be assigned to instances of a MIB object.

SNMP Community An SNMP community is an administrative relationship between an SNMP agent and one or more SNMP managers. Each community consists of a community name, an object access specification and a list of SNMP managers' IP addresses. A community is used by an SNMP agent to determine which requests are to be honored.

Heterogeneous Network A heterogeneous network is that in which a collection of systems of different type and manufacturer are interconnected by a variety of communication methods and protocols.

Request For Comments (RFC) A Request for Comments (RFC) is a technical report that documents standards, protocols, and guidelines for the development of TCP/IP protocol standards. RFCs are the mechanism by which TCP/IP and the Internet Protocol Suite are evolving. Research ideas and new protocols are documented and brought to the attention of the Internet community in the form of an RFC. Some RFCs describe protocols and applications that are so useful that they are recommended to be implemented in all future implementations of TCP/IP; that is, they become recommended protocols or de facto standards.

Request/Response Protocol A request/response protocol is one where in a communications environment the exchange of information among different entities is done through requests which are received by an entity for processing, after which it generates a response to be sent back to the originator of the request. SNMP uses this type of protocol to transfer data between managers and agents. The SNMP manager can send a request to the SNMP agent which will in return send a response.

SNMP Trap An SNMP trap is a message that is originated by an agent application to alert a managing application of the occurrence of an extraordinary event. SNMP traps include: coldStart, warmStart, linkDown, linkUp, authenticationFailure, EGPNeighborLoss, and enterpriseSpecific.

Object Identifier (OID) An object identifier is a means for identifying some object, regardless of the semantics associated with the object. An example would be a network object or a standards document. An object identifier is defined by ASN.1.

9.2.3 The SNMP Architecture

The SNMP architectural model is a collection of network management stations and network elements, such as gateways, routers, bridges and hosts. These elements act as servers and contain management agents which perform the network management functions requested by the network management stations. The network management stations act as clients; they run the management applications which monitor and control network elements.

SNMP provides a means of communicating between the network management stations and the agents in the network elements to send and receive information about network resources. This information can be status information, counters, identifiers, and more.

The SNMP manager polls the agents for error and statistical data. The performance of the network will be dependent upon what the polling interval is set at. The physical and logical characteristics of network objects make up a collection of information called a management information base (MIB). The individual pieces of information that comprise a MIB are called MIB objects, and they reside on the agent system. These objects can be accessed and changed by the agent at the manager′s request.

Unsolicited data, called traps, can also be sent from the agent to the manager under certain conditions. This is how NetView for OS/2 manages network objects. Other SNMP managers could also access these MIB objects.

9.2.4 Goals of the SNMP Architecture

The SNMP architecture explicitly minimizes the number and complexity of management functions realized by the management agent itself. This goal is attractive in that, among other benefits, it allows for the following:

- Reduced costs in developing management agent software to support the protocol
- Few restrictions on the form and complexity of management tools
- Simplified, easier to implement management functions

A second goal of the protocol is that the functionality can be extended to accommodate additional, possibly unanticipated, aspects of network management. A third goal is that the architecture be, as much as possible, independent of the architecture and mechanisms of particular hosts or gateways.

9.2.5 SNMP Model

The SNMP model is made up of the following components:

- At least one network element to be managed (agent system) containing an agent
- At least one network managing station (NMS), containing one or more network management applications
- A network management protocol for use by the NMS and the agent system to exchange network management information
- At least one MIB defining the information to be managed on the agent system

Figure 175 on page 393 is a graphical representation of the SNMP model.

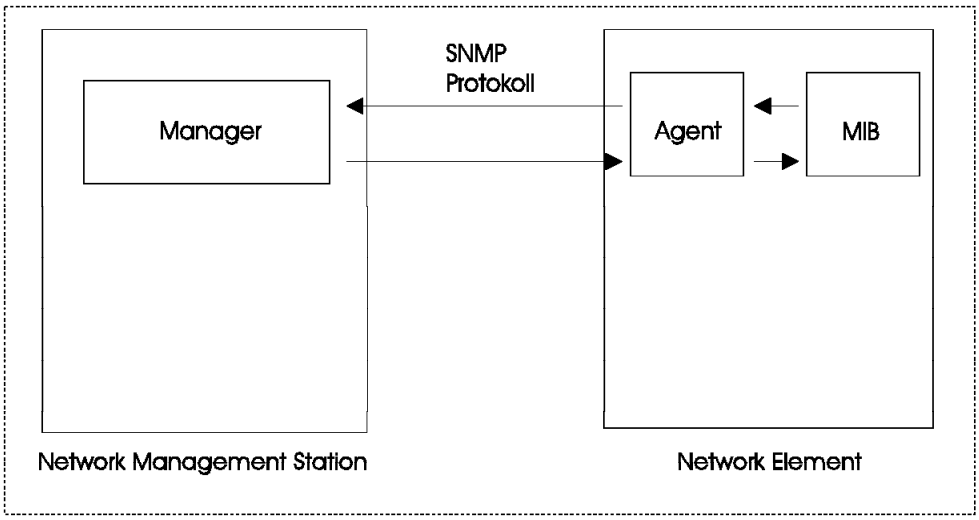


Figure 175. The SNMP Model

9.2.6 User Datagram Protocol (UDP)

The communication of management information among management entities is done in SNMP through the exchange of protocol messages, each of which is entirely and independently represented within a single UDP datagram using the Basic Encoding Rules (BER) of ASN.1. These protocol messages are referred to as protocol data units (PDU).

Consistent with the goal of minimizing complexity of the management agent, the exchange of SNMP messages requires a simple datagram service. For this reason, the preferred transport service for SNMP is the User Datagram Protocol (UDP), although the mechanisms of SNMP are generally suitable for use with a wide variety of transport services.

As a transport layer protocol, UDP uses the Internet Protocol (IP) as the underlying protocol. Two inherent characteristics of UDP provide for its simplicity. One of them is that UDP is unreliable, meaning that the UDP does not guarantee that messages will not be lost, duplicated, delayed, or sent in a different order. UDP is also a connectionless protocol, because the only process involved is the transfer of data. However, UDP does provide a certain level of data integrity validation through checksum operations. UDP also provides application layer addressing because it has the ability to route messages to multiple destinations within a given host. Figure 176 on page 394 shows where SNMP and UDP operate within the TCP/IP protocol stack.

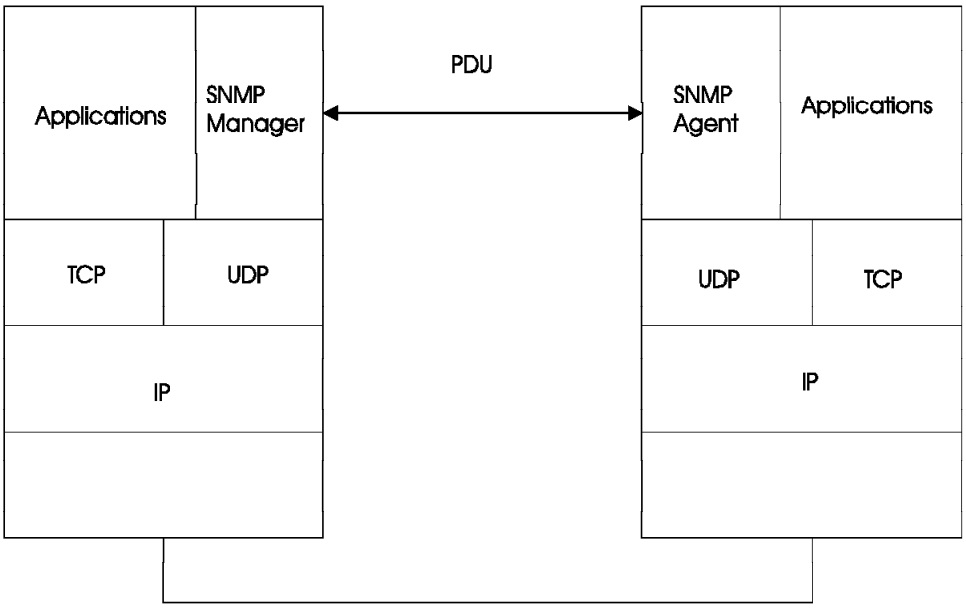


Figure 176. SNMP in the TCP/IP Protocol Stack

9.2.7 Asynchronous Request/Response Protocol

Managing systems generate SNMP requests, and agent systems generate responses to these requests. After a request message has been sent, SNMP does not need to wait for a response. SNMP can send other messages or realize other activities. These attributes make SNMP an asynchronous request/response protocol.

An agent system can also generate SNMP messages called traps without a prior request from the managing system. The purpose of a trap message is to inform the managing system of an extraordinary event that has occurred at the agent system. It must be noted that all request/response transactions are subject to the time delays inherent to all networks. The typical SNMP request/response primitives take place in the following manner:

- The manager polls agent with a request for information.
- The agent supplies information, which is defined in a MIB, in the form of a response.

Figure 177 on page 395 illustrates two time sequence diagrams. The top diagram shows a typical SNMP request/response interaction, while the bottom diagram shows a typical SNMP trap sequence.

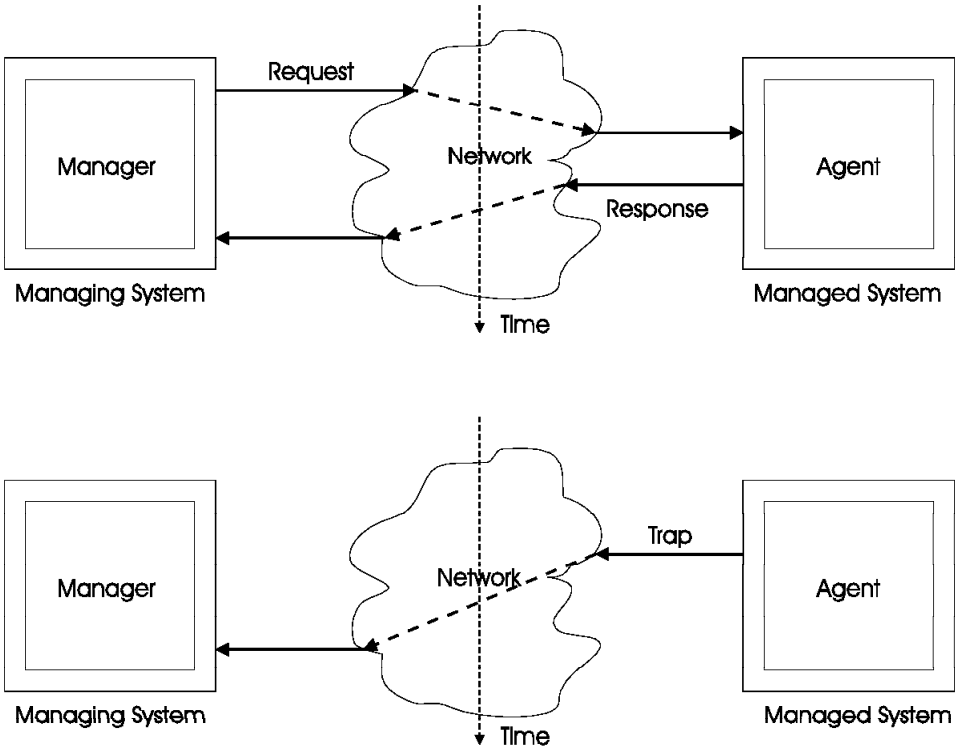


Figure 177. Asynchronous Request/Response Protocol

9.2.8 SNMP Agent

The SNMP agent has the following two responsibilities:

1. To gather error and statistical data defined by MIB objects.
2. To react to changes in certain MIB variables made by a managing application.

In summary, the following steps describe the interactions that take place in an SNMP managed network:

- The SNMP agent gathers vital information about its respective device and networks.
- The SNMP manager polls each agent for MIB information and can display this information at the SNMP manager station. In this manner, a network administrator can manage the network from a management station.
- An agent also has the ability to send unsolicited data to the SNMP manager in the form of a trap. A trap is generally a network condition detected by an SNMP agent that requires immediate attention by the network administrator.

9.2.9 SNMP Subagent

A subagent extends the set of MIB objects provided by an SNMP agent.

With a subagent it is possible to define MIB variables that are useful and specific to a particular environment, then register them with the SNMP agent.

Requests for the variable(s) that are received by the SNMP agent are passed to the process acting as a subagent. The subagent then returns an appropriate answer to the SNMP agent. The SNMP agent eventually sends an SNMP response with the answer back to the network managing station that initiated the request. The network management station has no knowledge that the SNMP agent calls on other processes to obtain an answer. From the viewpoint of the managing application, the agent is the only network management application on the managed system.

9.2.10 SNMP Manager

An SNMP manager refers to a network management station which runs a network management protocol and network management applications. SNMP is the network management protocol which provides the mechanism for management. Several different network management applications exist that can be used, such as NetView for OS/2, and NetView for AIX. The network management application provides the policy to be used for management.

The network management applications rely on management information base (MIB) objects for information regarding the managed system, also called the agent system. Management systems generate requests for this MIB information and an SNMP agent on the managed system responds to these requests. A request can either be the retrieval or modification of a MIB variable.

The agent system makes network and system information available to other systems by accessing the MIB objects and allowing configuration, performance, and problem management data to be managed by the SNMP manager.

For example, a network manager can access the system description of a particular agent system by using the network management application to gain access to the agent system′s sysDescr MIB object. To do this, the managing application builds a message that requests a MIB object called sysDescr. This request is sent to the agent system where the agent decodes the message and then retrieves the information related to the sysDescr MIB object. The agent constructs a response with this information and sends it back to the managing application. When the application has decoded the response, the SNMP manager can then display the agent system′s description information to the user. Figure 178 on page 397 shows the relationships among the SNMP entities as discussed in the previous paragraphs.

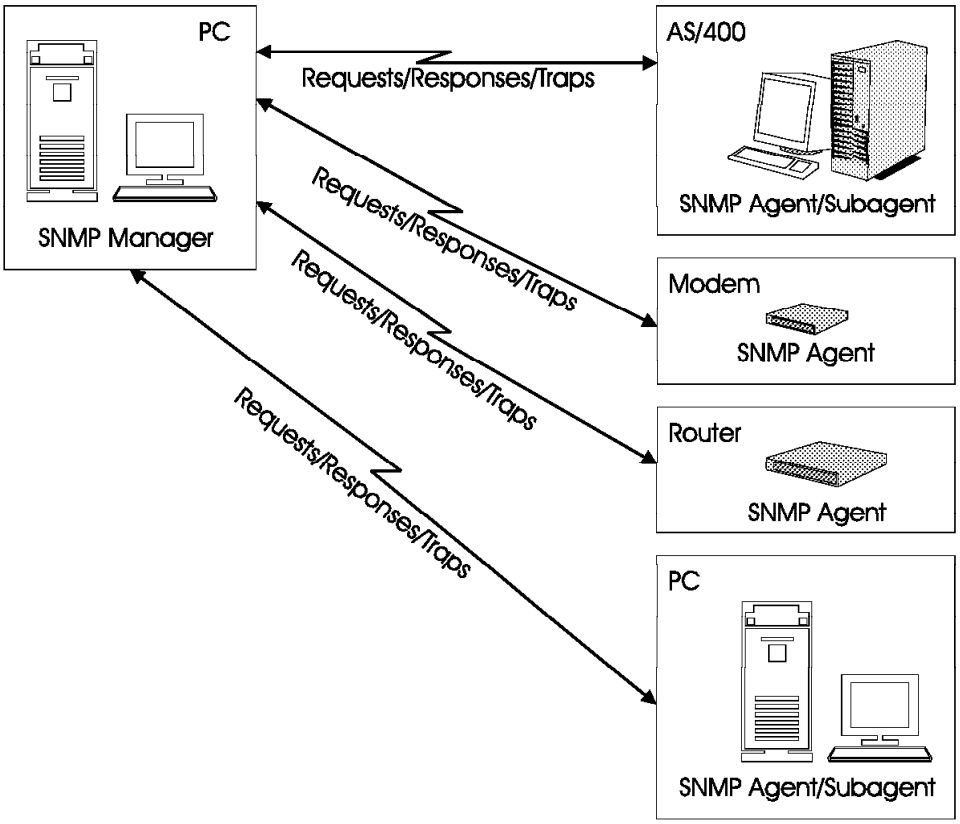


Figure 178. SNMP Manager/Agent/Subagent Relationship

9.2.11 SNMP Version 2

SNMPv2 is a new version of SNMP; it is documented in twelve RFCs. SNMPv2 was developed in order to give a better response to security and operational problems.

Up-to-date SNMPv2 information can be obtained by accessing the following World Wide Web site:

```
http://www.snmp.com/v2star.html
```

9.2.11.1 Security

In the original SNMP, the administrative relationship between an agent and one or more management applications was identified by a community. The community relationship involved the following three aspects:

- Identification of the entities authorized to request management operations
- Identification of the type of management operation that is allowed (read, write or none)
- Identification of management information that is available to the operations (MIB views)

Now with SNMPv2, three new concepts appear:

- The party concept which is an execution environment residing in an agent or management application, which refers to entities that

communicate via a management protocol and a transport service using authentication and privacy facilities.

- The context concept refers to collections of managed objects resources accessible by an SNMPv2 entity.
- The access policy concept defines the operations that may be performed when a source party communicates with a destination party and references a particular context. There are three levels of authentication/protection:

 snmpPrivMsg contains the party name and an snmpAuthMsg the content of which is encrypted by secret key.

 snmpAuthMsg contains authentication credentials and information about the management operation and its execution environment.

 snmpMgtCom contains the name of the party that originated the message, the party that is intended to receive the message, the managed objects, and the desired operation.

9.2.11.2 Operational Model

Some of the operations of SNMP remained the same and some were added. The following is a list of the operations available in SNMPv2:

- GET: This operation experienced no change.
- GETNEXT: This operation experienced no change.
- SET: This operation experienced no change.
- GETBULK: This operation was introduced to minimize network interactions, by allowing the agent to return large packets. This operation gets everything under the MIB. The number of variables that should be retrieved (non-repeaters) and the maximum number of times that other variables should be retrieved (max-repetitions) can be specified. If non-repeaters is greater than or equal to the number of variables in the request or non-repeaters is equal to zero and max-repetitions equal to one, a GETNEXT operation would be being emulated.
- INFORM: This operation is used when a management application wishes to inform another management application of some information. This operation always receives a response from the other management application.

9.2.11.3 SNMPv2 RFCs

The new SNMPv2 framework is defined in the following twelve RFCs:

- RFC1441 Introduction to SNMPv2
- RFC1442 SMI for SNMPv2
- RFC1443 Textual Conventions for SNMPv2
- RFC1444 Conformance Statements for SNMPv2
- RFC1445 Administrative Model for SNMPv2
- RFC1446 Security Protocols for SNMPv2
- RFC1447 Party MIB for SNMPv2
- RFC1448 Protocol Operations for SNMPv2

- RFC1449 Transport Mappings for SNMPv2
- RFC1450 MIB for SNMPv2
- RFC1451 Manager-to-Manager MIB
- RFC1452 Coexistence between SNMPv1 and SNMPv2

For more information on how to request RFCs refer to A.1.1, “Request for Comments (RFC)” on page 560

9.2.12 Understanding MIBs

The physical and logical characteristics of a system make up a collection of information which can be managed through SNMP. The individual pieces of information make up MIB objects. A Management Information Base (MIB) is comprised of MIB objects that reside on the agent system, where they can be accessed and changed by the agent at the manager′s request.

The administrative policy established by the IAB, results in a classification of MIBs according to their applicability and purpose. Therefore, MIBs are classified as follows:

Standard MIB: All devices that support SNMP are also required to support a standard set of common managed object definitions of which a MIB is composed. The standard MIB object definition, MIB-II, enables you to monitor and control SNMP managed devices.

Experimental MIB: Generally, new ideas and activities related to the Internet result in the definition of MIB objects. An experimental MIB is comprised of such objects. This approach offers the advantage that all new ideas must be proven while under experiment before they can be proposed for standardization.

Enterprise-Specific MIB: SNMP permits vendors to define MIB extensions or enterprise-specific MIBs, specifically for controlling their products. An enterprise-specific MIB must follow certain definition standards just as other MIBs must, to ensure that the information they contain can be accessed and modified by SNMP agents.

9.2.13 SNMP Operations

To be consistent with its simplicity objective, SNMP contains few commands to execute its operations. SNMP supports two commands that managing systems can use to retrieve information from a managed system and one command to store a value into a managed system. All other operations are considered to be side-effects of these three commands.

As an example, SNMP does not contain an explicit reboot command. However, this action might be invoked by simply setting a parameter indicating the number of seconds until system reboot. In addition to these commands, SNMP supports an event-driven mechanism used to alert managing stations of the occurrence of extraordinary events at a managed system.

The approach that SNMP is based on for network management makes it a simple, stable, and flexible protocol because it can accommodate new operations as side-effects of the same SNMP commands acting upon new MIB variables; thus not requiring SNMP to be modified.

SNMP also specifies that if a single SNMP message specifies operations on multiple variables, the operations will either be performed on all variables or on none of them. No operation will be performed if any of the variables are in error.

Each SNMP operation is defined in a particular PDU, a brief description of each operation follows.

- **GET.** This is a request originated by a managing application to retrieve an instance of one or more MIB objects. The specified instance is retrieved for each variable in the request, provided that community profile authentication has been successful.
- **GETNEXT.** This is a request originated by a managing application to retrieve the next valid instance following the specified instance of one or more MIB objects, provided that community profile authentication has been successful.
- **SET.** This is a request originated by a managing application to store a specific value for one or more MIB variables. All variables must be updated simultaneously, or none of them.
- **GET-RESPONSE.** This is response data that is originated by an agent application and is sent back to the originator of a GET, GETNEXT, or SET request.
- **TRAP.** This is an unsolicited message originated by an agent application which is sent to one or more managing systems within the correct community, to alert them of the occurrence of an event. Traps include the following types:
 - coldStart (0)
 - warmStart (1)
 - linkDown (2)
 - linkUp (3)
 - authenticationFailure (4)
 - egpNeighborLoss (5)
 - enterpriseSpecific (6)

9.2.14 Desktop Management Interface (DMI)

Within a computer, there is a gap between management software and the system′s components that require management. Managers must understand how to manipulate information on a constantly growing number of products. In order for products to be manageable, they must know the intricacies of complex encoding mechanisms and foreign registration schemes. This arrangement is not desirable from either side.

Therefore the Desktop Management Taskforce designed the Desktop Management Interface, or DMI, that acts as a layer of abstraction between these two worlds.

The DMI has been designed to be:

- Independent of a specific computer or operating system
- Independent of a specific management protocol

- Easy for vendors to adopt
- Usable locally, no network required
- Usable remotely using DCE/RPC, ONC/RPC, or TI/RPC
- Mappable to existing management protocols (for example, SNMP)

The DMI procedural interfaces are specifically designed to be remotely accessible through the use of remote procedure calls. The RPCs supported by the DMI include:

- DCE/RPC
- ONC/RPC
- TI/RPC

The DMI has four elements:

1. A format for describing management information
2. A service provider entity
3. Two sets of APIs, one set for service providers and management application to interact, and the other for service providers and components to interact
4. A set of services for facilitating remote communication

Component descriptions are defined in a language called the Management Information Format, or MIF. Each component has a MIF file to describe its manageable characteristics. When a component is initially installed into the system, the MIF is added to the (implementation-dependent) MIF database.

DMI Service Providers expose a set of entry points that are callable by component instrumentation. These are collectively termed the Service Provider API for Components. Likewise, component instrumentation codes expose a set of entry points that are callable by the DMI Service Provider. These are collectively termed the Component Provider API. In the DMI Version 1.x specifications, these two APIs were together embodied in the Component Interface.

The Component Interface, or CI, is used by component providers to describe access to management information and to enable a component to be managed. The CI and MIF shield vendors from the complexity of encoding styles and management registration information. They do not need to learn the details of the popular and emerging management protocols.

The DMI Service Provider also exposes a set of entry points callable by management applications. These are collectively termed the Service Provider API for Management Applications. Likewise, management applications expose a set of entry points callable by the DMI Service Providers. These are collectively termed the Management Provider API. In the DMI Version 1.x specifications these were together embodied in the Management Interface.

The Management Interface, or MI, is used by applications that wish to manage components. The MI shields management applications vendors from the different mechanism used to obtain management information from elements within a computer system.

For more information about the DMTF and DMI see http://www.dmtf.org.

9.3 Overview of IBM Products for Network Management

In this section we give you an overview about the IBM products in this area and about the different management platforms. For further information about the functions and interoperability of the products see *Network Operations Management - Which Platform? The Principles,* SG24-501 4 and *Network Operations Management - Which Platform? The Practice,* SG24-5015.

To be able to compare the different management platforms, we distinguish the following three different IT environments:

- LAN Workgroup

 This environment comprises PCs connected by LANs, where the LAN supports a group of people (for example, in a department). The typical IT resources found in LAN workgroup environments are:

 - PC-based file servers (for example, Novel NetWare and IBM LAN Server)
 - PC desktops that access file server resources (for example DOS, Windows and OS/2)
 - LAN bridges and hubs

- Distributed

 This environment consists of multiple LANs connected to each other, to form a dispersed internetwork. The typical IT resources found in distributed environments are:

 - PC-based file servers
 - UNIX systems
 - Mid-range systems (for example, DEC and AS/400)
 - PC desktops
 - LAN bridges and hubs
 - Routers

- Centralized

 This environment consists of multiple LANs and WANs connected to a host, where the host acts as a centralized server and data repository. Centralized environments include the IT resources found in distributed environments plus:

 - Mainframe systems
 - Communication controllers (for example, the IBM 3745 controller)
 - Switches (for example, ATM switches)

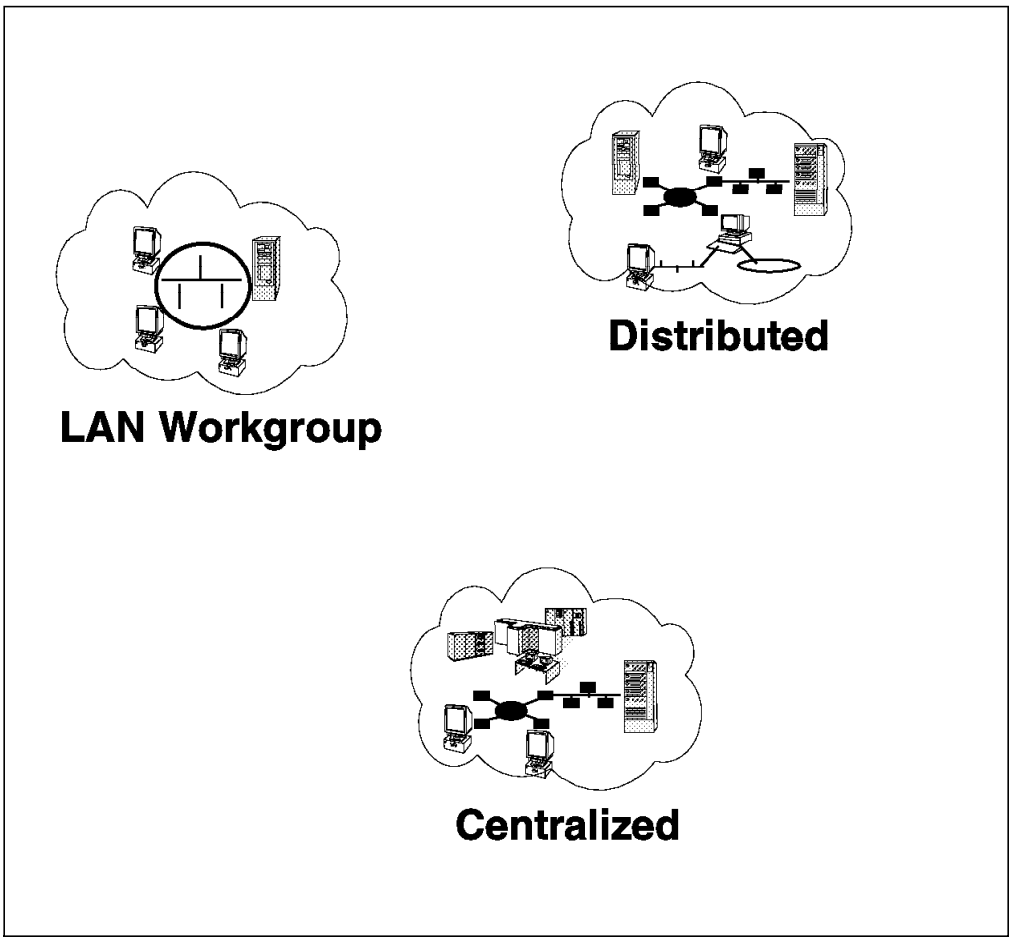


Figure 179. Categorization of IT Environments

9.3.1 Positioning the AIX Management Platform

The AIX management platform is a suitable candidate to manage distributed environments with heterogeneous, multi-vendor resources connected to a TCP/IP network. It also supports non-IP environments. The AIX management platform can manage thousands of devices, and it supports very dynamic networks with high rates of topology change. It supports requirements for high availability of the enterprise management system.

The environments where AIX may be a potential candidate management platform are:

- LAN Workgroup

 This environment consists of PCs connected by LANs, where the LAN supports a group of people. These PC LANs typically include file servers (for example, Novell NetWare and IBM LAN Server), PC desktops (for example, DOS, Windows and OS/2), bridges and hubs.

- Distributed

 This environment consists of multiple LANs connected to each other to form an internetwork. These internetworks typically include file servers, UNIX systems, mid-range systems (for example, DEC and AS/400), PC desktops, bridges, hubs and routers.

You should consider the AIX management platform if you require support of open industry standards, such as SNMP. AIX is an open platform with several interfaces (for example, the SNMP API) for application integration. The AIX management platform offers many applications from multiple vendors to manage an open, heterogeneous environment. Today there are about 130 applications for this platform.

The primary strength of the AIX management platform is managing IP networks using SNMP, but it can also support non-IP environments (for example, PCs in NetBIOS LANs) because it interoperates with multiple intermediate managers.

You may consider using the AIX management platform in an SNA, MVS-based environment where there are a growing number of IP devices. The AIX management platform interfaces with the MVS management platform in an SNA network.

The AIX management platform requires UNIX, TCP/IP and LAN skills to set up and maintain its multiple products.

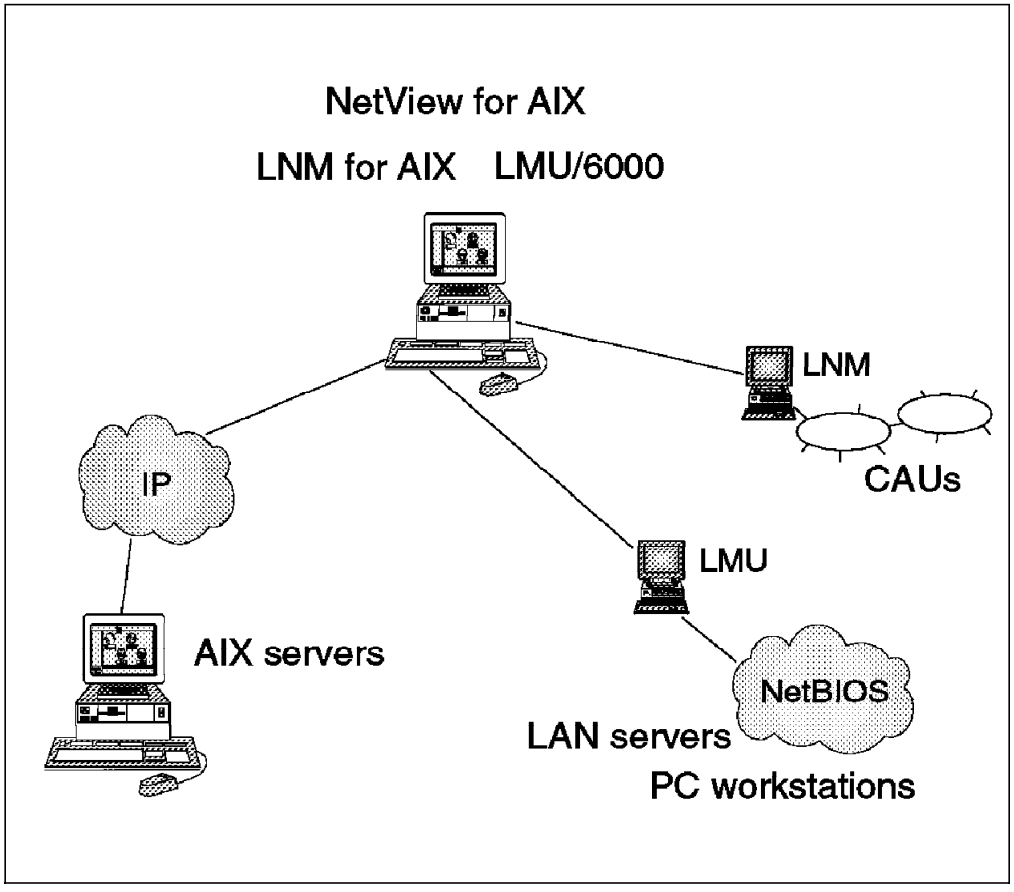


Figure 180. Example of an AIX Managed Network

9.3.2 AIX Management Platform Overview

The AIX management platform is an SNMP platform for managing heterogeneous network devices and systems in distributed environments. The main product is NetView for AIX, which manages IP networks, SNMP devices and other non-IP resources. NetView for AIX interoperates with the OS/2 and MVS management platforms to support cooperative management across the enterprise.

The AIX management platform can scale up to support thousands of devices. It can manage larger environments distributed across remote locations using UNIX-based, mid-level managers. These mid-level managers manage IP networks locally, relieving the load on the wide area network and NetView for AIX.

The AIX management platform can maintain high availability of the managing system. NetView for AIX has manager backup capabilities. When one NetView for AIX manager fails, another can take over and monitor its managed environment.

NetView for AIX interfaces with other intermediate managers to support non-IP environments. These intermediate managers run proxy agents that natively manage the non-IP networks. Two of these proxy agents are products from the OS/2 management platform:

- LAN Network Manager for OS/2 (token-ring LANs)
- LAN NetView Management Utilities (NetBIOS and IPX PC LANs)

The AIX management platform interfaces with NetView for MVS with two products: the AIX NetView Service Point and the SNA Manager/6000. The AIX NetView Service Point enables centralized management of IP networks from a focal point MVS platform. The SNA Manager/6000 manages SNA subarea networks from NetView for AIX (it requires NetView for MVS as the underlying SNA management engine). The number of SNA resources that can be managed with SNA Manager/6000 is limited.

9.3.2.1 Current Product Releases

The AIX management platform is well suited for heterogeneous multiprotocol environments. It interoperates with the OS/2 and MVS management platforms. The AIX management platform is comprised of AIX operating system features and several systems management products. The way the products fit together is described in *Network Operations Management - Which Platform? The Practice,* SG24-5015.

The products here were up-to-date for all general announcements made in most countries up to the end of May 1996:

- NetView for AIX V4 including the Openmon PTF
- LAN Management Utilities/6000 V1
- SNA Manager/6000 V1.1
- Router and Bridge Manager/6000 V1.2
- LAN Network Manager for AIX V1.0
- LAN Remote Monitor for AIX V1
- Nways Campus Manager ATM for AIX V2.1

- Nways Campus Manager LAN for AIX V2.1
- Nways Campus Manager for AIX V2
- Nways BroadBand Switch Manager R3
- Telecommunications Management Network Product Family for AIX
- Trouble Ticket for AIX V3.2
- Systems Monitor for AIX V2
- AIX NetView Service Point
- Various products from the NetView Association

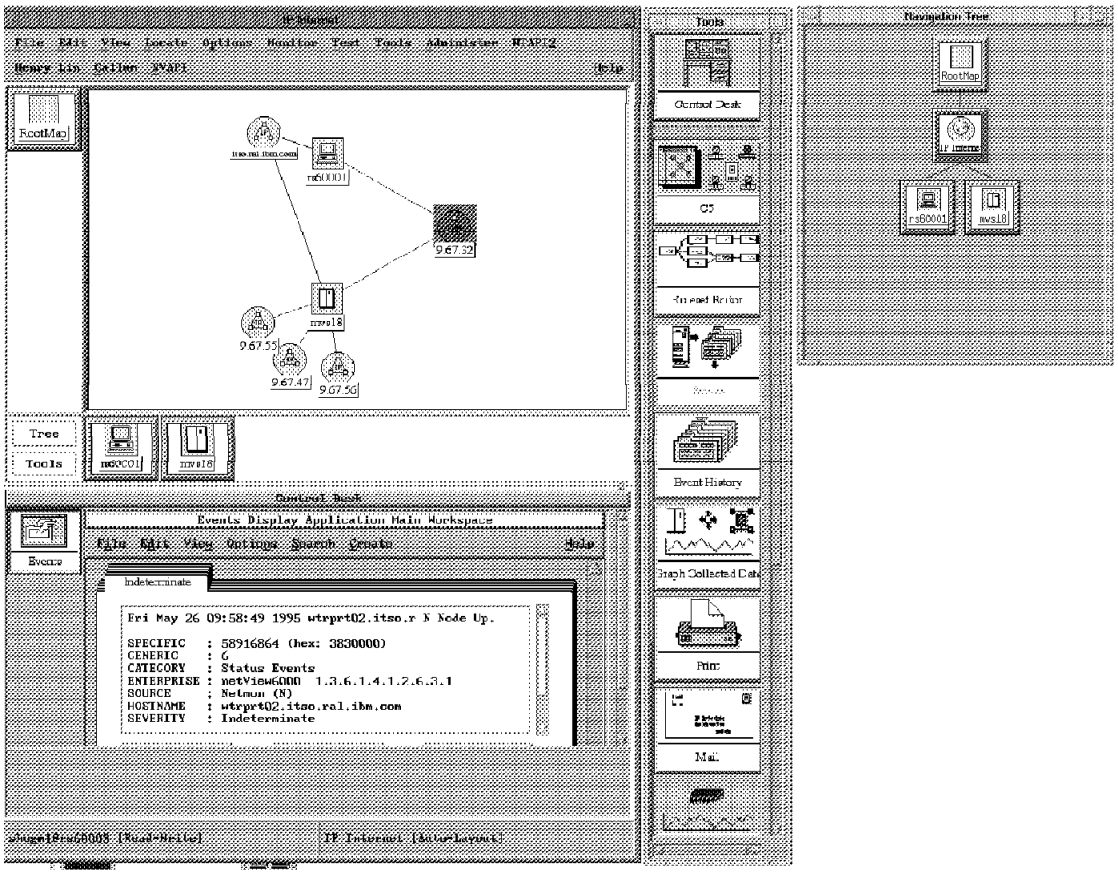


Figure 181. The NetView for AIX Desktop including Navigation Tree and Tool Palette

9.3.3 Positioning the MVS Management Platform

The MVS management platform is a suitable candidate to manage centralized (mainframe-centric) and distributed multi protocol environments connected to an SNA network. The MVS management platform can manage tens of thousands of devices and it supports very dynamic networks with high rates of topology change.

If you do not have MVS in your environment, it would probably not make sense to consider MVS as a candidate for management platform. If you already have MVS, the MVS platform may be a potential candidate to manage practically all types of environments:

- LAN Workgroup

 This environment consists of PCs connected by LANs, where the LAN supports a group of people. These PC LANs typically include file servers (for example, Novell NetWare and IBM LAN Server), PC desktops (for example, DOS, Windows and OS/2), bridges and hubs.

- Distributed

 This environment consists of multiple LANs connected to each other to form an internetwork. These internetworks typically include file servers, UNIX systems, mid-range systems (for example, DEC and AS/400), PC desktops, bridges, hubs and routers.

- Centralized

 This environment consists of multiple local and wide area networks connected to a mainframe. Centralized environments include mainframe systems, communication controllers, switches (for example, ATM switches), OEM equipment and all the resources found in distributed environments.

You should consider the MVS management platform a very strong candidate if you have MVS, and use an SNA network; the MVS management platform allows you to leverage your staff′s MVS skills. You should also consider MVS as a candidate management platform if you require very high availability and reliability in your environment. MVS is the most mature and stable of the IBM management platforms.

The MVS management platform′s primary strength is managing large SNA networks, but it can also support large heterogeneous environments because it interoperates with multiple intermediate managers. Some of the non-SNA environments supported by the MVS management platform are:

- NetWare LANs
- IP networks
- Token-ring LANs

You may consider the MVS management platform if you have an existing SNA, MVS-based environment and you also require support of open industry standards (for example, DCE and SNMP).

The MVS management platform offers sophisticated functions, but it requires extensive skills (MVS, SNA, etc.) to set up and maintain its environment and the multiple products that run on it.

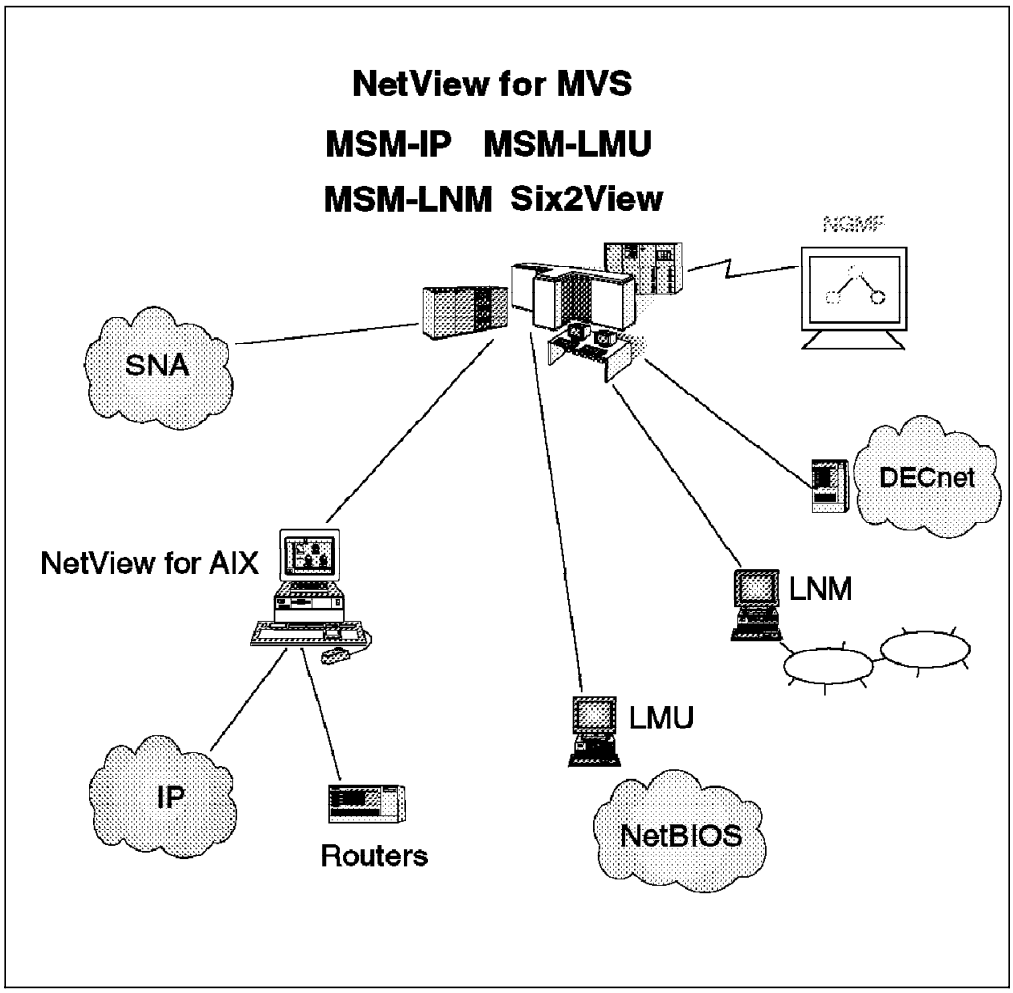


Figure 182. Example for an MVS Solution

9.3.4 MVS Management Platform Overview

The MVS management platform allows you to centrally manage distributed and mainframe-centric environments from one focal point manager. The main product is NetView for MVS, which can manage SNA networks and other non-SNA environments. The MVS platform provides very sophisticated functions for systems and network management, including extensive automation support.

The MVS management platform can manage large multi protocol environments because it interoperates with multiple intermediate managers. The intermediate managers run service point applications that natively manage the distributed environments. Some examples of these service point managers are:

- NetView for AIX and the AIX NetView Service Point (IP environments)
- NetWare for SAA and the NetWare Management Agent for NetView
- LAN Network Manager for OS/2 (token-ring LANs)
- LAN NetView Management Utilities (NetBIOS and IPX PC LANs)

NetView for MVS managers can cooperate with each other on a peer-to-peer basis. You can assign different spheres of control to different NetView for

MVS managers, and use one of them as your enterprise-wide focal point manager.

NetView for MVS implements an object-oriented, in-memory repository of data about managed resources. This data cache is called the Resource Object Data Manager (RODM). The RODM object-oriented infrastructure enables multiple applications to share managed resource information and use it to integrate and automate their functions.

The MVS management platform provides centralized operations, problem, configuration, performance, change and business management. It can closely integrate these systems management processes because the NetView for MVS platform offers many interfaces that have been exploited by multiple products from the NetView for MVS family (for example, NetView Performance Monitor) and other MVS-based systems management applications (for example, Information Management).

9.3.4.1 Current Product Releases

This list shows the release levels we used when we wrote the following section. The subset of these products needed in a given enterprise depends on its complexity and how much integration you want with other processes. See *Network Operations Management - Which Platform? The Practice,* SG24-5015 for more information about the usage of the products in different environments.

Program products:

- NetView for MVS V2.4 including:
 - For monitoring of SNA:
 - NetView Graphic Monitor Facility (NGMF)
 - NetView APPN Topology and Accounting Management Feature (APPNTAM)
 - To monitor any non-SNA:
 - NetView Resource Object Data Manager (RODM)
 - NetView Graphic Monitor Facility Host Subsystem (GMFHS)
 - NetView MultiSystems Manager V2.2 (MSM) including
 - The OS/2 LAN Network Manager Networks Feature
 - The Novell NetWare Networks Feature
 - The LAN NetView Management Utility Networks Feature
 - The TCP/IP Networks Feature
 - To add intermediate managers:
 - AIX NetView Service Point V1R2
 - NetView for OS/2
 - For managing AS/400s:
 - NetView Remote Operations Manager MVS
 - NetView Remote Operations Agent /400
 - To manage digital equipment:

- Six2View from Phoenix Network Technologies Inc.

– For automating the link to problem management:
 - NetView AutoBridge/MVS V1R1
 - Information Management V6R2
– To automate performance management:
 - NetView Performance Monitor V2R2 (NPM) including:
 • NPM Desk/2
– For configuration management:
 - NetView Network Planner/2 V1R2 (NNP/2)
– Miscellaneous:
 - Open Systems Interconnection Communication Subsystem (OSI/CS) V2

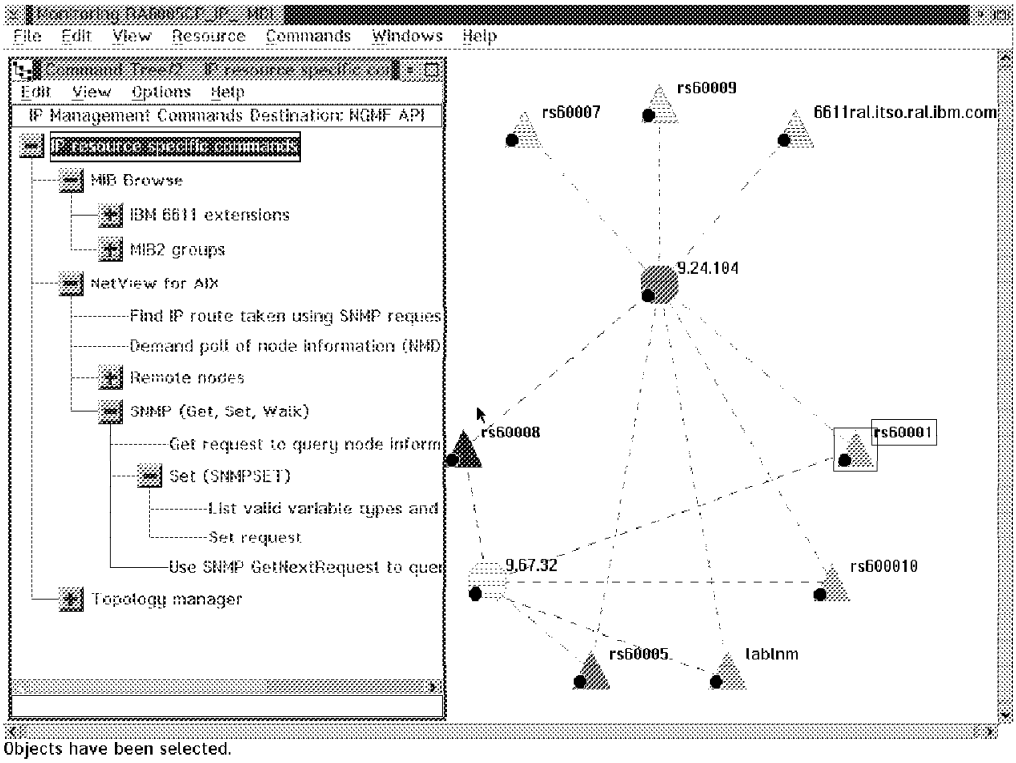


Figure 183. MSM View of IP Resources

9.3.5 Positioning the OS/2 Management Platform

The OS/2 management platform is a suitable candidate to manage LAN environments with PC systems and multiple network protocols (TCP/IP, NetBIOS, IPX and SNA). The OS/2 management platform can manage hundreds of devices, and it focuses on PC systems, instead of on network devices. It offers limited support for dynamic networks with changing topology (only token-ring LANs). The OS/2 management platform interoperates with the AIX and MVS management platforms to support cooperative management across the enterprise.

The primary environment where OS/2 is a potential candidate management platform is the LAN Workgroup. This environment consists of PCs connected by LANs, where the LAN supports a group of people. These PC LANs typically include file servers (for example, Novell NetWare and IBM LAN Server), PC desktops (for example, DOS, Windows and OS/2), bridges and hubs.

The OS/2 management platform supports some of the complexity found in a distributed environment, because it offers limited support for hubs, routers, UNIX and mid-range systems. The distributed environment consists of multiple LANs connected to each other to form an internetwork. These internetworks typically include file servers, UNIX systems, mid-range systems (for example, DEC and AS/400), PC desktops, bridges, hubs and routers.

You should consider the OS/2 management platform if you require support for the SNMP industry standard, but do not have UNIX in your LAN environment.

The OS/2 management platform requires OS/2, LAN and PC skills to set up and maintain its environment. Its skill requirements are greater than those for the Windows platform.

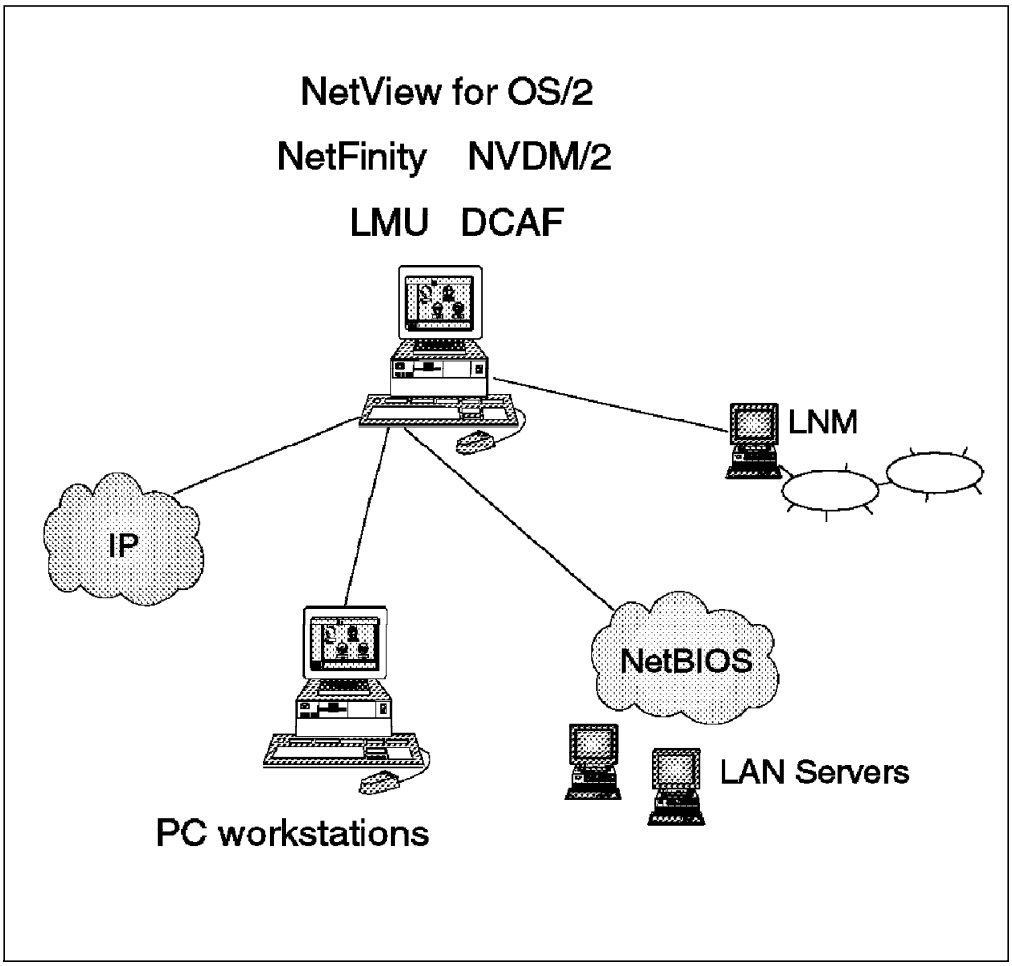


Figure 184. OS/2 Managed Network

9.3.6 OS/2 Management Platform Overview

The main product of the OS/2 management platform is NetView for OS/2, which manages SNMP devices in TCP/IP, NetBIOS, IPX and SNA networks. NetView for OS/2 is a low-cost SNMP management platform compared to UNIX-based SNMP managers, which run on more expensive RISC platforms. NetView for OS/2 is an open platform with several interfaces (for example, the SNMP API) for application integration.

NetView for OS/2 can scale up to support hundreds of devices. It can manage medium-size environments distributed across remote locations using intermediate LMU (LAN NetView Management Utilities) managers. These LMU managers manage NetBIOS (IBM LAN Server) and IPX (NetWare) networks locally, off-loading the backbone network and NetView for OS/2.

The OS/2 management platform interfaces with NetView for MVS to enable centralized management of SNA-connected PC LANs from a focal point MVS platform. It also interfaces with NetView for AIX to enable centralized management of IP-connected PC LANs from NetView for AIX.

9.3.6.1 Current Product Releases

The OS/2 management platform is comprised of OS/2 operating system features and several systems management products:

- NetView for OS/2 V2.1
- LAN NetView Management Utilities (LMU) V1.1
- IBM SystemView Manager for OS/2 V1.1
- LAN Network Manager for OS/2 (LNM) V2.0
- System Performance Monitor/2 (SPM/2) V2.0
- Distributed Console Access Facility (DCAF) V1.3
- Network Door/2 (NetDoor) V1.0

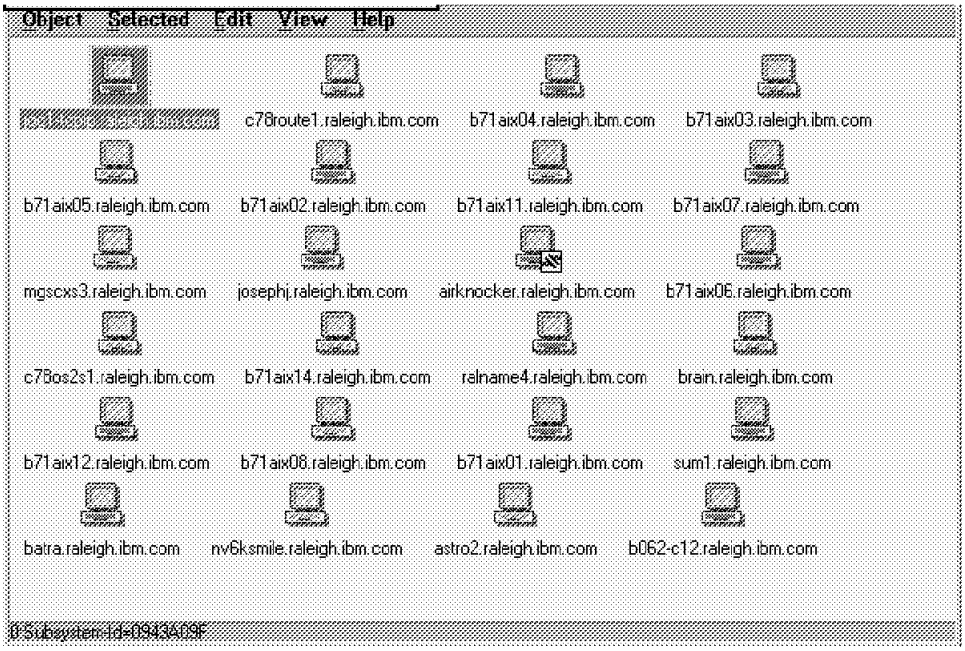


Figure 185. NetView for OS2 All Resources Status Display

9.3.7 Positioning the Windows IBM Management Platform

The Windows 3.1 IBM management platform is a suitable candidate to manage small LAN environments with heterogeneous network devices (bridges, hubs and routers) and PC systems. The Windows IBM management platform can manage around a hundred resources and it focuses on network devices and ease of use. It does not support dynamic networks with a lot of topology changes.

The Windows IBM management platform does not interoperate with the AIX, MVS or OS/2 management platforms. Therefore, it does not fit in an enterprise environment.

The only environments where the Windows IBM management platform is a potential candidate are small LAN Workgroups. This environment consists of PCs connected by LANs, where the LAN supports a group of people. These PC LANs typically include file servers (for example, Novell NetWare and IBM LAN Server), PC desktops (for example, DOS, Windows and OS/2), bridges and hubs.

You should consider the Windows IBM management platform if you require SNMP support, but do not have UNIX or OS/2 in your LAN environment.

The Windows IBM management platform requires Windows, LAN and PC skills to set up and maintain its environment. Since this platform is easy to set up and use, it does not require extensive skills.

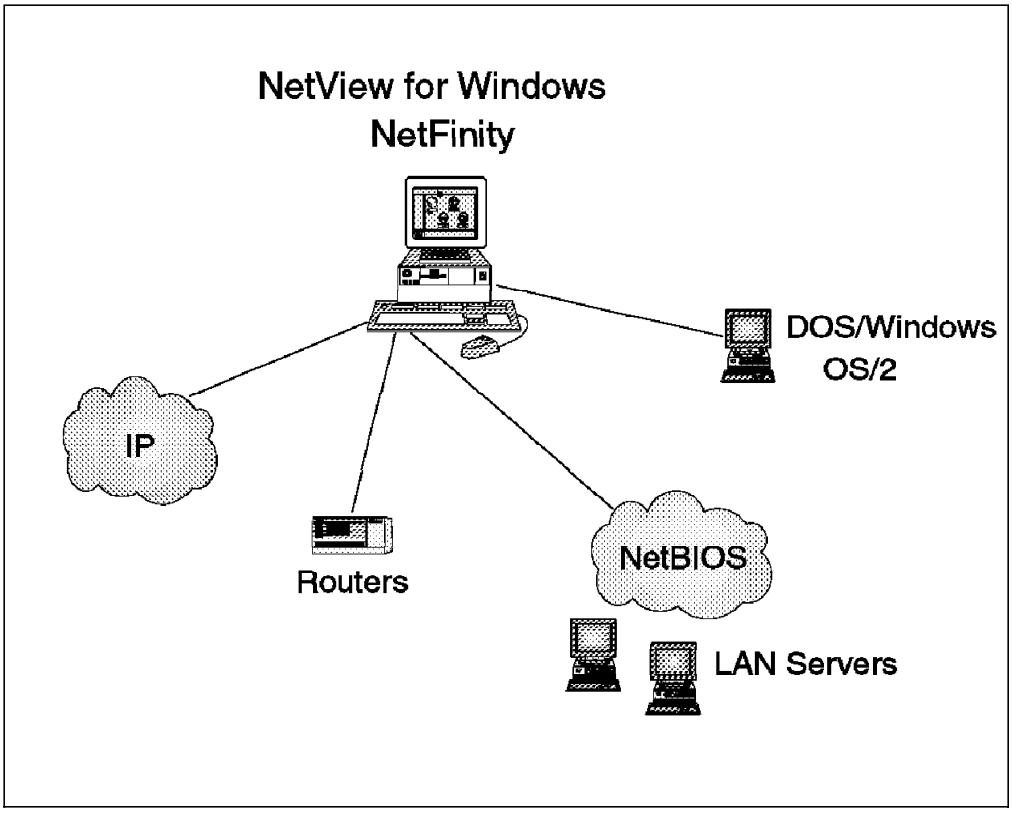


Figure 186. Network Management with Windows Manager

9.3.8 Windows IBM Management Platform Overview

The key product of the Windows IBM management platform is NetView for Windows, which manages SNMP network devices in TCP/IP LANs. NetView for Windows is a low-cost SNMP management platform compared to UNIX-based SNMP managers, which run on more expensive RISC platforms. It focuses on managing network devices (bridges, hubs and routers), with limited support for PC systems.

NetView for Windows supports many network management applications that are provided by different vendors, to manage their network devices. There are two types of these device-specific applications, which provide the following two levels of function:

- Basic management applications, also known as PIMs or product integrator modules.
- Advanced management applications, also known as PSMs or product specific modules.

The following are the reasons why the Windows IBM management platform fits only small LANs:

- NetView for Windows network maps must be customized manually, and are not updated for dynamic topology changes. This is acceptable only for small LANs with a moderate rate of change.
- The Windows platform does not offer the reliability required to manage medium or large LANs.
- The Windows IBM management platform meets the requirement (in small LANs) that ease of use should have a higher priority than function.

9.3.8.1 Current Product Releases

The Windows IBM management platform assessed in this chapter is comprised of these products:

- NetView for Windows V2.0
- NetFinity Manager for Windows V3.06
- LAN Remote Monitor for Windows V1
- Nways Manager for Windows

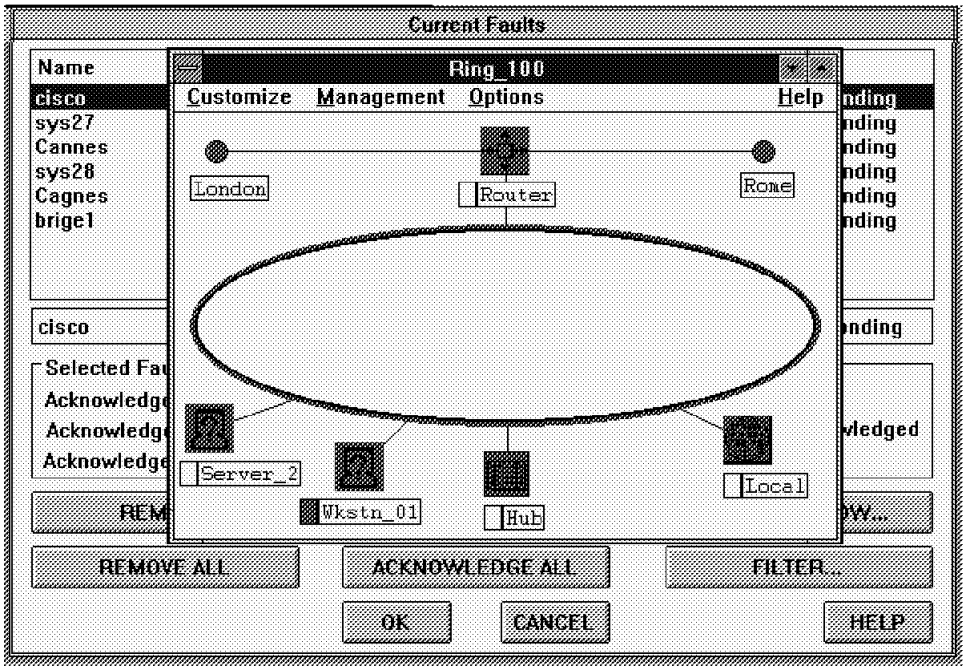


Figure 187. Topology Overview with NetView for Windows

9.3.9 Tivoli TME 10

TME 10 is IBM′s management solution for client/server, enterprise management, the Internet, and beyond.

With TME 10, you can standardize on a care set of systems management functions across the enterprise, picking among the best-of-breed tools to put together the total solution.

TME 10 is a winning combination that consists of Tivoli TME solutions, IBM SystemView solutions and industry solutions (hardware and software vendors, database and application vendors, and other systems management vendors). The primary driving force behind the Tivoli and IBM merger is the synergy (the complementary strengths) that exist between the companies and their products. The cross-platform products from both product lines are represented in TME 10. Consolidating the product lines was not a process of choosing one offering over another. Instead, it centered on defining points of integration and selecting the best-of-breed features that existed in each product category.

TME 10 is based on a single architecture and object-oriented framework (the Tivoli Management Framework) which is based on open standards to enable its common applications and third-party applications to run on a diverse set of management platforms. One of the primary benefits of an object-oriented framework for systems management is integrating a variety of complimentary management applications without having to re-write the entire application. This strategy allows you to use management tools created by different organizations together as an integrated whole.

The customer gains scalability by defining what functions are needed where and who is allowed to run them. Platform independence is achieved in that the various difference between the supported management platforms are

hidden from the operator or administrator. Customers can pick and choose where they want those functions to reside. The management console integrates systems, network, and applications management together from a single place.

TME 10 supports the following hardware platforms, with appropriate operating systems support:

- IBM RISC System/6000 and PowerPC Systems
- NCR (formerley AT&T) System 3000
- Data General AViiON systems
- Motorola 88000 series systems
- Sun SPARC systems
- Intel 486 or Pentium, or equivalent
- HP 9000 systems

The complete roll-out of TME 10 will occur in three general phases:

1. Packaging

 Consolidate product offerings in each functional area of TME 10. Create single orderable products where several alternatives exist. Clearly identify those cases where a particular product will be phased out, and a migration path to the preferred TME 10 product.

2. Application Integration

 Create a single, integrated product offering in each functional area. Previously separate products integrated to form a single, cooperative management product.

3. Framework Integration

 Migrate all underlying services onto a common framework; eliminate any overlapping management console interfaces.

9.4 More Information

If you need more information about SystemView, SNMP or IBM products for managing heterogeneous networks, see *SystemView for AIX V1R1: Scenarios,* SG24-2564, *The Basics of IP Network Design,* SG24-2580, *Network Operations Management - Which Platform? The Principles,* SG24-5014 and *Network Operations Management - Which Platform? The Practice,* SG24-5015.

Internet user can get information about redbooks and IBM products from the following URLs:

- http://www.redbooks.ibm.com/redbooks
- http://www.software.ibm.com
- http://www.raleigh.ibm.com/nethome.html
- http://www.software.ibm.com/sysman/
- http://www.tivoli.com

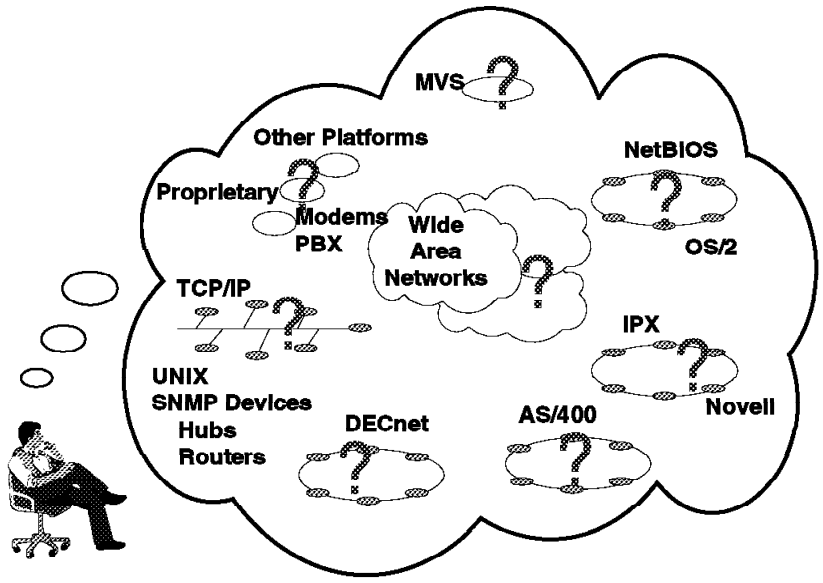


Figure 188. Which Platform Should Be Used As the Manager?

Chapter 10. Connection Access Services

This chapter describes the IBM Internet Connection Access Services. Before that, we make an explanation about what a service provider is, IBM as a server provider, how to select one, and how to build an infrastructure for an Internet Service Provider (ISP).

For additional information, refer to:

- http://www.ibm.com/globalnetwork/inetcnbr.htm
- http://www.ibm.com/globalnetwork/cb9502.htm

10.1 Service Providers

A service provider is a company that has a dedicated Internet gateway which is shared by companies and individual users. Some providers have more than a dedicated gateway to the Internet; they have a backbone network. Many people already have access to the Internet through a service provider and don't even know it. Your company may provide corporate access into the Internet through a corporate gateway. Some of them just provide mail access. To access the Internet properly, you need a TCP/IP network connection.

10.1.1 How to Select an Internet Service Provider

Buying an Internet connection is a lot like buying a computer. Just like when you are buying a computer, your choice of an Internet service provider should be driven by your intended use. If you are looking for a minimum cost, you might seek out the lowest-priced system in the back of a magazine or even assemble something yourself from parts bought at a flea market. However, if you are buying something for your company that your business will depend on, you would probably make different choices.

For your business, you might consider buying the most expensive solution, exercising the theory that you get what you pay for. However, once you have really studied this question, the right choice might well turn out to be a mid-range system from a stable, nationally recognized provider.

There are some low-cost IP service suppliers who claim to be just as good as the others, but may not be in business next year to prove it. Also, there are other suppliers who will attempt to justify providing the same level of services as their competitors, at many times the price.

Some questions we need to think about to evaluate service providers are the following:

- ***Network Topology***: Network topology is one of the most important criteria to consider when choosing a provider. Looking at the network topology can help you understand how vulnerable the network is to outages, how much capacity is available when the network is loaded more heavily than usual, and, the most important, how well the provider understands network engineering.

© Copyright IBM Corp. 1996

Any competent service provider should be happy to show you their network topology. This is a good way for them to demonstrate how well they understand their business.

Look closely at what they show you, some providers will give you a virtual backbone map. Virtual networks are meaningless. Your data does not flow on a virtual network but it flows on a physical one. A virtual network map is merely a representation of all the theoretical paths that could be implemented by the provider's virtual circuit switching equipment and it is an attempt to side-step the issue of physical capability. Your supplier needs to understand the physical network to understand what is important for serving their customers. If they tell you that the physical topology is unimportant, they either don't understand how to engineer a network or they are trying to disguise something. It is important to say that there is nothing inherently wrong with using frame relay, or other technologies that use virtual circuits as part of the backbone. However, your provider must understand the physical topology on top of which their virtual (logical) network is running.

- ***Network Link Speeds***: It is important to look closely at the speeds of the backbone links. If they won't show you these speeds, then they are probably hiding something. The first thing to understand is that your network connection can only be as fast as the slowest link in the path.

 It doesn't matter if you are connected to a T-3 node if there is a 56 kbps link between you and your destination. The limit is the 56 kbps link, not how much capacity the T-3 node has.

 Next, ask if the topology you are being shown is operational now. Some providers like to show links that are not operational as part of their backbone infrastructure. It is also important not to be confused between the press release about a new high-speed network link and that link actually being operational.

- ***External Network Links***: Take a look at the external links of each provider's backbone. Do they have a single connection to the rest of the world? This is a potential single point of failure. Look for multiple, direct connections to other network providers. The more of these connections, the better. This shows that the provider is concerned about external connectivity and does not want to be dependent on some third party for interconnection. If they have a single connection to the outside world, ask them how often it fails and how long they usually are isolated. If they can't give you these statistics, are they managing their own network well enough to manage yours?

- ***High-Speed Backbone***: If they claim to have a high-speed backbone, check to see if it is that speed now or if it is just planned. Some providers claim to have a T-3 (45 Mbps) backbone, but if pushed, will admit that what they really mean is that it is T-3 capable.

 The next thing to ask yourself about high-speed backbones is if you can actually connect to it for a reasonable cost. All service providers require you to buy the local loop segment from your facility to their closest point of presence, or POP. You will have to buy this directly or indirectly from one of the telephone companies serving your local area. Some providers offer their service in such a way that the local loop cost is greater than their fee to provide you with the service in the first place. If you're limited by the local loop speed because the price of a high-speed loop is not cost effective, then how useful is a high-speed backbone?

- ***Technology***: The technology being used to operate the network is also critically important. Today, there is plenty of commercial quality router, switch and modem technology available from companies whose business is to make that equipment. Any provider still relying on their own internally developed equipment is doing you a disservice. You deserve the benefits of leading-edge production technology, not aging hardware that has been contorted into a use never intended by its designers.

 Sometimes a provider can have a bad case of the *not invented here* syndrome. This is a sure sign of long-term problems. Remember, you are buying a service. The provider of this service should be using the best available technology to deliver this service.

- ***Build or buy?***: Some providers claim that they need to run even the lowest layers of their network to deliver quality service. This is not true. The truth of the matter is all Internet service providers rely on one or more telephone companies to assemble their network. The only way for any company to build their own network is to physically dig their own trenches and lay their own fiber into the ground.

 The only real question is at which physical link or transport level your potential service provider buys from the much larger phone companies. If the lower-level infrastructure and service (such as T-1, T-3, frame relay or ATM) needed to support an Internet service provider′s value added service is offered by a phone company, it′s not cost effective or in the best interest of the provider′s customers for the provider to even think about building and operating it. The provider simply cannot match the economy of scale that comes with being a phone company. If your provider has chosen to build something when they could have bought a more reliable service more cheaply, why should you have to pay for their misplaced priorities? The job of an Internet service provider is to manage and maintain its IP level connectivity.

- ***Technical Staff***: One of the most important aspects to consider when choosing a provider is the quality of their technical staff. They are the ones who will get your connection running to begin with and then keep it and the network running in future. They have to be experienced in TCP/IP data networking.

 Make sure the provider has adequate staffing to cover the usual situations. If they send people to trade shows for a week, how many people are back at the office running things and how skilled are they? Find out what their technical staff turnover is. If people are leaving, find out why and who is left to keep your connection operational. Many suppliers of service have single points of failure in their staff capacity as well.

- ***Help Desk Infrastructure***: Check out their help desk infrastructure. It should be 7x24 (24 hours a day and 7 days a week) staffed by at least one person, including nights, weekends, holidays and during important sport events. Make sure that they will have someone capable of dealing with your problem and not someone who will just answer the phone all the time.

- ***Organization***: Find out how long the company has been in the IP business. Try to determine if they are going to be in business for the long run. Quality networks are not built on a little budget. The pricing may look attractive now, but the passage of time often reveals hidden costs and price increases, the greatest of which can be having to switch

providers. Ask about their financial stability, if they have a positive cash flow and are going to be in business next month to provide your connectivity. Determine if they have one or two major accounts that provide a disproportionate amount of revenue and what impact losing those accounts would have on their ability to keep its quality of service.

- ***Full Range of Services***: Does your provider have a full range of services or is it just filling a niche? If you need to increase or decrease your service level, will you need to switch providers?

 Does your provider offer true one-stop shopping? Can they supply equipment, manuals, training, consulting, etc., as well as basic services? Can they provide connectivity throughout the country and the rest of the world or do they just serve a small region? Can they provide service in other countries through established partnerships with international suppliers and bill you on the same invoice as your domestic service?

- ***Price/Benefit Analysis***: Do a price/benefit analysis. Some providers may appear to be priced less than others. Make sure you do an apples-to-apples comparison. Don't compare one one-service provider with another's full-service offering. Don't be confused by the names of the products. What one provider thinks is basic may be useless to you.

- ***Conclusion***: The amazing, worldwide growth of the Internet as a public access computer network has all kinds of new users, large and small, investigating the virtues of getting on the Internet. Today, more and more companies are using the Internet to conduct their business, communicate with and support their customers, exchange electronic mail with hundreds of thousands of users, and seek and find valuable information leading to competitive advantage. This resource is indispensable once obtained. The choice of the service provider to be responsible for ensuring this vital business tool is the most important decision you will make when embarking on the Internet.

10.1.2 How to Build an Infrastructure for an Internet Service Provider

This section describes what is needed to build an infrastructure for an Internet Service Provider (ISP) from in a corporate LAN.

An ISP has to connect its corporate systems up to an IP router and a leased line to the Internet. To access the Internet properly, you need a TCP/IP network connection and you can have a leased line connection to IGN and have full access to all sites in the Internet. You also can be a direct gateway into the Internet. As an ISP, you will be able to decide which services you will offer to your customers or corporate users.

When setting up a corporate link into the Internet, you need to take a number of things into account. These include:

- What speed of communication is required?

 The speed of this link will be driven by the number of users you plan to provide this service to and also the number of applications and data types that you will be using. Most ISP use either a 56 or 64 kbps line.

 It's hardly recommended that you give special attention to the increase of customers and corporate users to have the basis to plan the link upgrade.

- What line options do you have?

Line options include E1, T1, ISDN, and analog 56 kbps and 64 kbps (see Table 31 on page 423).

Table 31. Line Options

Service Grade	Speed	Notes
Standard Voice	0 to 28.8 kbps	SLIP, PPP, or dial-up connections.
ISDN	56 or 64 to 128 kbps	Digital phone line required; worldwide availability sporadic; common in Europe; dedicated or dial-up.
Leased	56, 64, 128, 256 or 512 kbps	Dedicated link to a service provider. Full TCP/IP access.
T1	1.544 Mbps	Dedicated link with heavy use.
T2	6 Mbps	Not commonly used in networking.
T3	45 Mbps	Major networking artery for a large corporation or university.

- How are you going to manage your security?

 Your corporation will have a full access connection to the Internet. Along with this access comes a large problem: security. Although the corporation now has access to the Internet, your corporate LAN will be opened to access from the Internet. Your corporate users, customers and all the Internet users will have access to your corporate network. If this unrestricted access is not a problem for you (maybe it's important for your business that all Internet users have full access to the information in your corporate LAN), you don't have anything to worried about. But if you want to avoid this, you should install a firewall at the Internet connection point. With firewalls, a company can make selected data and applications accessible to the Internet, while sensitive data is restricted.

 Firewalls and Internet security are detailed in Chapter 8, "Security on the Internet" on page 339.

 Additional information about firewalls and Internet security, refer to:

 – *Building an Infrastructure for the Internet,* SG24-4824-00
 – *Building a Firewall with the IBM Internet Connection Secured Network Gateway*, SG24-2577-01
 – URL: `http://www.ics.raleigh.ibm.com`

- What Internet services do you want your customers and/or corporate LAN users to use?

 As an Internet Service Provider (ISP), you need to decide what Internet services will be available for your customers and/or corporate users. Based on your decision, you'll need to choose which application servers you'll install in your corporate LAN.

 Following are some application server types:

 – News server
 – FTP server
 – Gopher server
 – WWW server
 – SMTP and POP servers

Most Internet users start by using the system to send electronic mail. Mail involves sending an electronic mail message to a user@location.

SMTP (Simple Mail Transfer Protocol) is the underlying transmission mechanism for much of the Internet mail. SMTP is a simple peer-to-peer model. Each host that wants to receive mail will set up an SMTP server. When mail is sent, it will be received by the SMTP server. You will then contact the local SMTP server to look at your mail.

POP (Post Office Protocol) is a protocol designed to handle the problems of having to log into the mail server to get your mail, and rather than have customers bring up the mail from the mail server. The POP server must be running POP-compliant code. The customer will then contact the POP server which will transmit the customer's mail to the customer.

The latest POP version is POP V3, or POP3, Post Office Protocol 3.

– Proxy servers

 A proxy server, or application gateway, secures traffic for a particular TCP/IP application. The proxy server will authenticate users for remote applications. Proxy servers are normally used for security reasons, such as in a firewall.

– Socks servers

 A socks server intercepts and redirects TCP/IP requests that cross between two portions of the Internet. The socks server will intercept each TCP/IP request, validate its userID, and check for authorization to go into or out of one are of the network to another. Applications such as Telnet, FTP, Finger, Gopher, Mosaic and Web Explorer can be handled through a socks server. In such a way, a socks server can pass Internet traffic without the traffic violating the system security.

– Name server

 It's important for you, as an ISP, to show your customers that you are a direct gateway into the Internet (even though you may not be). You will have an Internet domain company_name.com, such as ibm.com, and your customers will have e-mail user IDs as userID@company_name.com. You will be able to have your own Web page available www.company_name.com so that people will be able to find out about your service.

 For additional information about Domain Name System, refer to:

 - Chapter 11, "Content Services on the Internet" on page 451
 - *Accessing the Internet*, SG24-2597-00

For additional information about Internet Services, refer to:

– *Using the Information Super Highway*, GG24-2499-00

– *Accessing the Internet*, SG24-2597-00

10.1.2.1 Network Solution Design

Figure 189 shows a sample network solution design for an Internet Service Provider (ISP). You can use IBM RISC/6000 and AIX or PowerPC as servers in this solution. IBM 2210, IBM 6611 or Cisco routers can be used to connect your corporate LAN to the Internet and the IBM 8235 DIALs providing LAN remote dial-in access. All of this hardware attachs directly to either an Ethernet or a token-ring LAN.

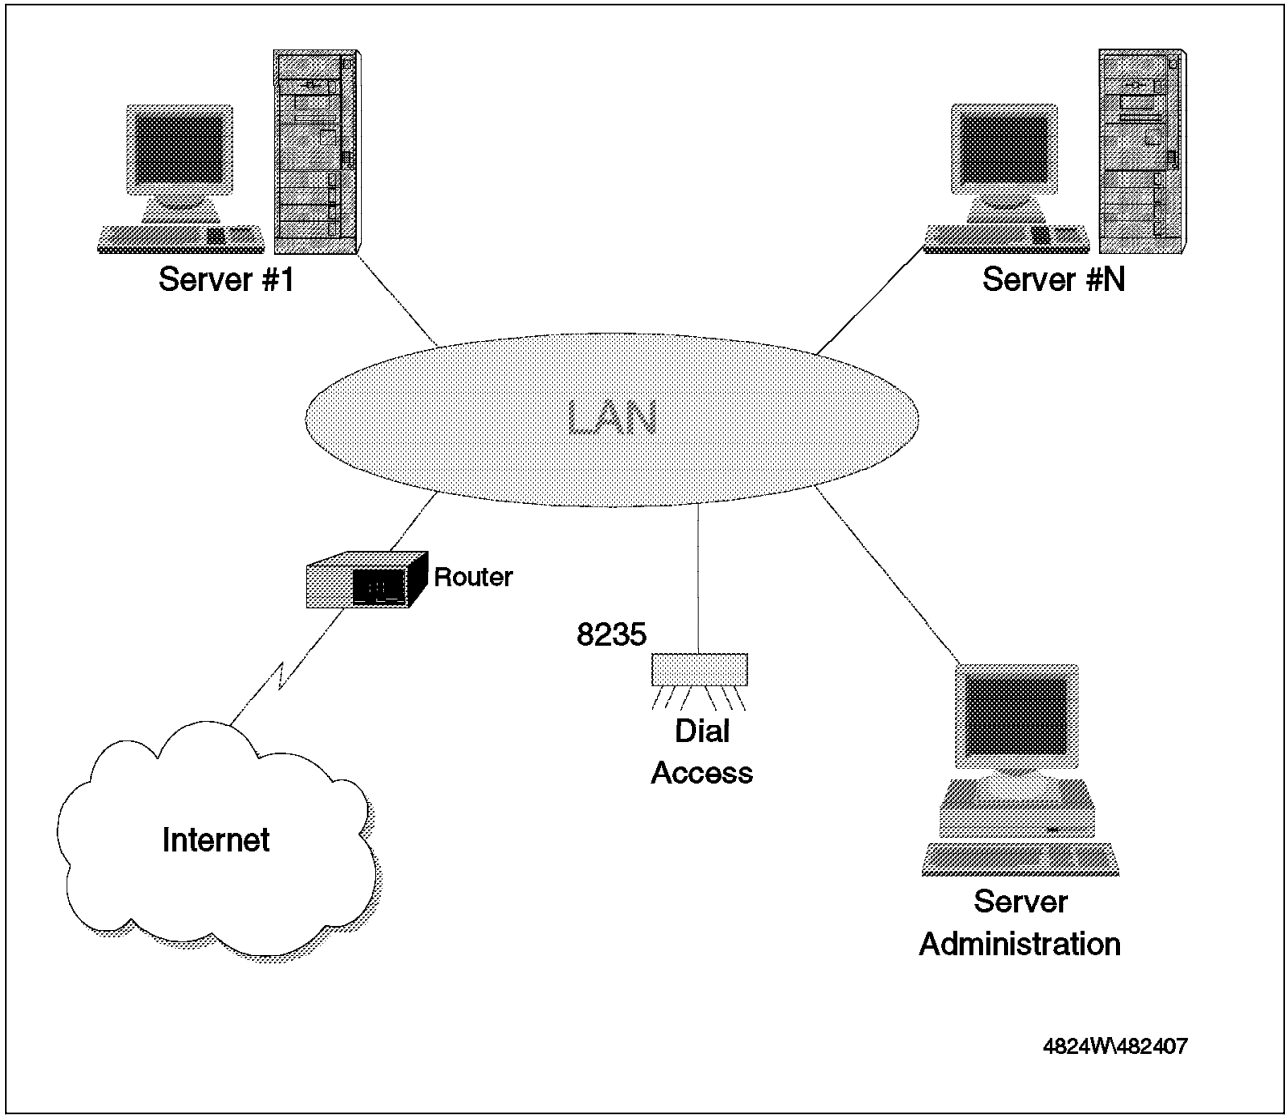


Figure 189. Proposed Network Solution Design for an Internet Service Provider

For detailed information about:

1. IBM RISC/6000, refer to:
 - http://www.austin.ibm.com/indext.html
2. PowerPC, refer to:
 - http://www.chips.ibm.com/products/ppc
3. IBM 2210 Nways Multiprotocol Router, refer to:
 - Chapter 2, “Networking Hardware” on page 21

- *Local Area Network Concepts and Products: Routers and Gateways*, SG24-4755-00
- *IBM 2210 Nways Multiprotocol Router Description and Configuration Scenarios*, SG24-4446-01

4. IBM 6611 Router, refer to:
 - Chapter 2, “Networking Hardware” on page 21
 - *Local Area Network Concepts and Products: Routers and Gateways*, SG24-4755-00
 - *MPNP V1R3 for IBM 6611*, SG24-4597-00
5. IBM 8235 DIALs.
 - Chapter 2, “Networking Hardware” on page 21
 - *IBM 8235 Dial-In Access to LANs Server - Concepts and Experiences*, SG24-4816-00

10.2 IBM As a Service Provider

IBM has set up networks and communication connections to service providers all around the world. These service provider connections have been combined with IBM′s vast network resources to form the IBM Global Network. This global network provides access to more than 90 countries and 700 cities. IBM provides different services for users accessing the Internet and offers the following service provider options:

- IBM Global Network
- Advantis network offerings
- Prodigy service offerings

Advantis and Prodigy are the largest IBM linked service providers in the USA. Both Advantis and Prodigy companies are joint ventures formed by IBM and Sears, Roebuck and Co.

Internally, IBM has access to the Internet through over 600 gateways in 50 countries at speeds up to 28.8 kbps via the IBM Global Network. IGN will offer dial access from 750 locations by year-end, and dial access speeds up to 64 and 128 kbps later this year via ISDN. IBM′s internal/external proxy and socks gateways are managed by tollbooth machines.

10.2.1 IBM Global Network

To provide international support for users wishing to access the Internet, IBM set up the IBM Global Network. This is a commercial service that provides end users with the advantage of IBM′s worldwide networking resources. IGN operates the world′s largest high-speed network for telecommunications services and network-centric computing. It brings together IBM′s capabilities to provide seamless, value-added network services globally through Advantis, the IBM Information Network organizations worldwide, and wholly-owned subsidiaries and joint ventures around the world. IGN has 5,000 network professionals and provides access to more than 90 countries and 700 cities. It provides value-added network services to more than 25,000

IBM customer accounts and, in many cases, to their vendors and suppliers. Additional worldwide advantages include:

- Local dial access numbers
- Low-cost Internet connectivity
- Leased-line access to the Internet
- Gopher, News and World Wide Web servers that assist you in navigating the Internet
- Worldwide customer support
- Integrated connectivity support with the OS/2 Warp operating system

For additional information, refer to:

```
http://www.ibm.com/globalnetwork
```

10.2.2 Advantis

Advantis is a network service provider, as it provides the physical connectivity on the Internet. End users can register with Advantis as their Internet connection provider. As such, Advantis is responsible for setting up all the high-speed network connections, SMTP and POP servers, domain name servers, routing, Internet IP administration, etc. Advantis forms part of the IBM Global Network as the USA and Canadian Internet service provider of the IBM Global Network service. Advantis provides SLIP and dedicated leased lines as connection options. The Advantis leased line connections range is from 56 kbps to 1.544 Mbps.

For additional information, refer to:

```
http://www.advantis.com
```

10.2.3 Prodigy Services Company

Prodigy Services Company is a consumer-oriented online information service company. It provides services over and above simple Internet information management. Internet users cannot access the Prodigy data directly. They must first sign onto a Prodigy account. Prodigy Services Company assimilates vast amounts of information gained from numerous sources and brings them together in a usable form. Prodigy offers its members a range of news, computing, weather and sport, financial information, educational content, games, reference materials, communications features such as e-mail, newsgroups and Chat, travel reservations, shopping, online banking, and other offerings.

The three major competitors in this area are:

- Prodigy
- Compuserve
- America Online

Prodigy users connect via a dial-up connection to a Prodigy server using dedicated Prodigy software. The user does not connect into the Internet and is not part of the Internet. While accessing the Prodigy system, the user can use a Web browser provided by Prodigy for accessing the Internet through a gateway. The user can send and receive e-mail on the Internet. A Prodigy

user cannot do a telnet or FTP into the Internet and, as such, is by no means a complete Internet user.

Prodigy uses Advantis as its link into the Internet.

For additional information, refer to:

`http://www.prodigy.com`

10.3 IBM Internet Connection Access Services

The IBM Global Network offers a secure, reliable and flexible set of high-speed, leased line Internet access solutions that can include network connectivity resources, and security options designed, installed and managed by the IBM Global Network. Customers can establish high-speed leased line access to the Internet, without having to install and manage their own network hardware, software and telecommunications links. They can choose the approach that best suits their requirements, one-way lanes to the Internet with firewall security options, or open direct access to the Internet over dedicated leased lines. The IBM Global Network also offers remote and mobile users access to the Internet via a local dial from over 600 points of presence around the world, and 24-hour, seven-day-a-week customer support.

10.3.1 Dial-Up Services

Dial access is provided via the IBM Global Network′s direct access backbone. Remote and mobile users may use a variety of software packages including IBM′s OS/2 Warp Internet Access Kit, IBM Internet Connection Access Kit, IBM Internet Connection for Windows, Netmanage′s Chameleon, and Ventana′s Internet Membership Kit.

10.3.1.1 Highlights

The IBM Internet Connection service is a comprehensive suit of access, applications and services to get customers on the road to the information superhighway.

Access

- Over 600 local dial access points for low-cost connectivity around the world
- 800 dial service for users outside local calling areas in the U.S. and Canada
- IBM′s dial service supports every major platform, including Windows, UNIX, Macintosh and OS/2. Dial users can choose one of the following commercial offerings for easy connection:
 - IBM Internet Connection access kit, including Netscape, Eudora Light E-Mail and Trumpet Winsock
 - IBM Internet Connection for Windows, including WebExplorer Mosaic, e-mail, NewsReader, Gopher, FTP and Telnet
 - IBM OS/2 Warp (Bonus Pack) and OS/2 Warp Connect, including IBM WebExplorer, Ultimedia Mail/2 Lite, NewsReader, Gopher, FTP and Telnet

 - Any SLIP protocol software can be used for IBM Global Network's dial access. IBM has set up sample scripts for popular software such as Trumpet Winsock, SPRY Internet in a Box, Netmanage Chameleon, Windows 95, MAC InterSLIP and LINUX (UNIX for PCs). PPP is currently not supported via our dial gateways. No dates are available at this time

 For additional information about these softwares and connection scripts, refer to:

 - http://www.ibm.net/software.html

- Up to 28.8 kbps connectivity for high-speed access (V.34 and V.42 support). IBM announced on June 18, 1996 that customers of the IBM Global Network will, by mid-July, be able to connect to IBM Global Network Internet dial service at increased speed. A new IBM platform that includes new modems introduced by U.S.Robotics will permit access at a speed of 33.6 kbps.

 IBM Global Network is able to accomplish enhancements like this quickly and easily through its new platform for dial services, called the local gateway interface, or LIG. The LIG, developed jointly by IBM and U.S.Robotics, features an IBM RS/6000 running AIX and a U.S.Robotics NAS (network access server) Chassis with modems or T1/E1 attachments. The LIG provides a common architecture for deploying IBM Global Network dial services, with many capabilities implemented in software.

 Advantis, the U.S. provider of the IBM Global Network, already uses the LIG platform in its Internet, TCP/IP and multiprotocol LAN dial services.

 Support for the new 33.6 kbps standard outside the U.S. is subject to availability of IBM Global Network dial services within a given country and will be rolled out in other geographies based on that availability. Today, the IGN platform for dial services outside the U.S. is called intelligent network gateway, or INGW. The INGW, developed by IBM, features an IBM PS/2 running OS/2 and a U.S.Robotics NAS Chassis with modems or a Motorola Codex Chassis with modems.

 For additional information:

 - about U.S.Robotics, refer to http://www.usr.com
 - about Motorola, refer to http://www.motorola.com

- Full TCP/IP connectivity with dynamic IP address assignment eliminating the need for customer to preregister an IP address.

 Direct dial access provides full TCP/IP connectivity via SLIP, along with support for all Internet protocols and applications, including Telnet, File Transfer Protocol (FTP), World Wide Web browsers, USENET/Newsgroups, SMTP e-mail, Gopher and Archie.

- The IBM Internet Dialer is used to establish a SLIP dial connection to the Internet through the IBM Global Network. Benefits of its use are:

 - Easy phone number selection and updates
 - Login assistance with error messages and retry
 - TCP/IP configuration assistance
 - Modem configuration assistance

 - Automatic server setup (name and mail servers, default Web page, etc.)
 - Pop-up messages at login to inform of new services and offerings
 - Online Internet registration and setup
 - Online and context-sensitive help and FAQs
 - Customer assistance links, such as: help desk phone numbers, automated connections to support newsgroups and Web pages, notify incident assistance, and e-mail problem reporting
 - Easy online updates of Dialer software
 - Easily configured application autostarting
 - Dial on demand support
 - Internationalization and NLS beyond most default dialers
 - Brandability for reselling and outsourcing
 - Connection logging and diagnostics
 - Application programming interfaces for third-party software
 - Automatic code updates
 - Inactivity timeouts and warnings
- Up to six user IDs available per subscription.

Applications

- POP3 (Post Office Protocol 3) servers available to hold your mail while you are not connected
- Up to 32-character mail names, for example, IBM_Corp_ITSO_redbooks_worldwide@ibm.net
- Change e-mail identity. The assigned ID, commonly known as the *user ID*, is used for both network access and e-mail access. They are limited in length and availability. This facility allows users of the IBM Internet Connection to choose a different e-mail ID which will offer more flexibility
- Convenient mail forwarding allows users of the IBM Internet Connection to forward mail to another Internet address
- Domain Name Server available to allow customers to use friendly, recognizable names when navigating Internet resources
- Default Gopher and World Wide Web (WWW) servers provided to help customers to begin their journey on the net
- News server which provides access in Internet news groups to follow different subjects, including discussion groups designed specifically for IBM Internet Connection users

Services

- 24-hour, seven-day-a-week customer assistance
- Local dial numbers for IBM Global Network′s Help Desk in almost every country IGN has a Internet point of presence
- Superior network management to provide timely access

- Usage details available online. This facility allows customer to obtain billing summary information regarding his account
- Major credit cards accepted
- Charges applied in local currency

10.3.1.2 Hardware and Software Requirements

- Hardware

 The recommended minimum hardware configuration for the IBM Internet Connection for Windows 3.1 and Windows 95, IBM Internet Connection Access Kit, and IBM Internet Connection for OS/2 is any personal computer with an Intel or 100% compatible 80386, or higher (recommended 80486, or higher) microprocessor, a minimum clock speed of 25 megahertz (MHz) and 8 MB of memory (RAM).

 Microsoft Windows 95 requires any personal computer with an Intel, or 100% compatible 80486, or higher, microprocessor and a minimum of 8 MB of memory (RAM).

- Software

 The IBM Internet Connection for Windows 3.1 and Windows 95, or the IBM Internet Connection Access Kit requires Microsoft Windows 3.1x, Microsoft Windows for Workgroups 3.1x, or Microsoft Windows 95. Also requires IBM Disk Operating System 5.0, or higher, or Microsoft Disk Operating System 5.0, or higher, and operates in Windows-enhanced mode. The IBM Internet Connection for OS/2 requires OS/2 Warp Version 3.0 or OS/2 Warp Connect.

- General system requirements

 The IBM Internet Connection for Windows 3.1 and Windows 95, the IBM Internet Connection Access Kit and the IBM Internet Connection for OS/2 require 15 MB of hard disk space, one 3.5-inch, 1.44 MB diskette drive, and a mouse, or compatible pointing device.

10.3.1.3 Connectivity

The IBM Internet Connection for Windows 3.1 and Windows 95, the IBM Internet Connection Access Kit and the IBM Internet Connection for OS/2 allow switched communication speeds up to 28.8 kbps. The effective speed will depend on the type of modem and serial port the modem is connected to. A Hayes-compatible modem supporting 9.6 kbps, or higher, and a telephone line are required. The following standards are currently supported:

- V.32 (9,600 bps)
- V.32bis (14,400 bps)
- V.34 (up to 28,800 bps)

The supported error control and data compression standards are:

- MNP Level 1-4 (error control)
- MNP Level 5 (data compression)
- V.42 (LAPM error control)
- V.42bis (data compression)

10.3.2 Corporate Dial Services

IBM Internet Connection corporate dial services is a dial offering by Advantis, the US provider of the IBM Global Network. This service allows corporate professionals, including workers in remote offices, telecommuters and business travelers, to access applications that reside on Transmission Control Protocol/ Internet Protocol (TCP/IP) hosts, servers and applications on the Internet.

In addition, users may access TCP/IP hosts and servers that are connected to the Advantis open IP network, including POP3 mail servers, news servers and WWW servers managed by Advantis. Users also have access to WWW content provided by Advantis, IBM and other companies. Connection is accomplished by placing a local phone call or an 800 call (if available and subject to surcharge) to one of the Advantis dial gateways on the Advantis high-speed IP network. The Advantis IP network is connected to the Internet at multiple network access points (NAPs), providing high-speed access to the Internet backbone.

10.3.2.1 Highlights

IBM Internet Connection corporate dial services provides dial access to the Internet using Serial Line Internet Protocol (SLIP) from personal computers or workstations with TCP/IP software. Users will have access to a full range of Internet applications and utilities such as NewsReaders, File Transfer Protocol (FTP), Archie, Gopher, Veronica, World Wide Web (WWW) and an optional offering for electronic mail. The billing for this service is handled through the standard Advantis billing process which produces invoice for these corporate customers. Internet applications and utilities are covered in detail in the redbook *Using the Information Super Highway*, GG24-2499-00.

IBM Internet Connection corporate dial services provides the following features:

- Local dial access from more than 350 cities in the US.
- Support for V.34 with dial access speeds up to 28.8 kilobits per second (kbps).
- Advantis provides a master copy of the IBM Internet Connection Corporate Access Kit for the Windows Version 1, Release 3.1 licensed software package as a part of the service. This package currently includes the Netscape Navigator WWW browser, Eudora Light Internet mail and Trumpet TCP/IP software in addition to the dialer.
- Users of IBM OS/2 Warp Version 3 may also use IBM Internet Connection corporate dial services through the dialer and Internet applications included in the BonusPak for IBM OS/2 Warp Version 3. In either case, registration for the service is accomplished by contacting an Advantis marketing specialist or IBM marketing representative. Lists of dial locations and modems are included in the respective packages and updates can be downloaded from the service. (See Appendix E, “IBM Global Network Phone List” on page 595 for the IBM Global Network Phone List.)
- Optional Internet mail accounts using Post Office Protocol 3 (POP3). If a company already has an Internet connection through Advantis or another Internet service provider, they have the choice of either maintaining an Internet mail post office on their server or using IBM Internet Connection

corporate dial services POP3 Internet mail accounts. The optional POP3 mail accounts include 5 megabytes (MB) of storage per user. The mail user ID is initially identical to the network access user ID assigned at the time of registration but the customer may change it to any available unique combination of up to 32 characters via a utility on the WWW.

- Optional customer-selectable custom mail domain name. If a customer chooses the optional mail service, the default mail domain name is ibm.net. For an additional one-time charge, a custom domain name may be used. If a company is already connected to the Internet through Advantis or a different Internet service provider and already has a domain name registered with the InterNIC, that domain name may be used. If a customer has not registered their domain name, Advantis will register their domain name choice with the InterNIC if that name is available, subject to InterNIC approval. In either case, the one-time charge for custom mail domain applies. This one-time charge for custom mail domain does not cover any InterNIC domain name registration or maintenance fees which will be billed directly to the customer by the InterNIC.
- User network authentication by account, user ID and password. Users of IBM Internet Connection corporate dial services connect to the Internet by first logging onto the Advantis network. Advantis provides the appropriate phone number, user IDs and initial passwords. Users may request passwords which expire every 60 days or less, or expire upon initial logon but are subsequently non-expiring.

 The user places a call to an Advantis dial gateway which authenticates the user′s account ID, user ID and password. This helps prevent unauthorized use of the Advantis network. Once the requester has been authenticated as a valid network user, the dial gateway assigns a dynamic IP address, sends it to the requesting device and the IP route to the Internet is established. At this point, the user can start one or more TCP/IP applications (for example, Telnet, FTP, NewsReader or WWW browser).
- Ability to use existing Advantis accounts, user IDs and passwords with this service.

10.3.2.2 Hardware and Software Requirements

- Hardware

 The recommended minimum hardware configuration for the IBM Internet Connection Corporate Access Kit for Windows 3.1 and Windows 95, and for the IBM Internet Connection for OS/2 is any personal computer with an Intel or 100% compatible 80386, or higher (recommended 80486, or higher) microprocessor, a minimum clock speed of 25 megahertz (MHz) and 8 MB of memory (RAM).
- Software

 The IBM Internet Connection Corporate Access Kit for Windows 3.1 and Windows 95 requires Microsoft Windows 3.1x, Microsoft Windows for Workgroups 3.1x, or Microsoft Windows 95. It also requires IBM Disk Operating System 5.0, or higher, or Microsoft Disk Operating System 5.0, or higher, and operates in Windows-enhanced mode. The IBM Internet Connection for OS/2 requires OS/2 Warp Version 3.0 or OS/2 Warp Connect.

- General system requirements

 The IBM Internet Connection Corporate Access Kit for Windows 3.1 and Windows 95, and the IBM Internet Connection for OS/2 require 15 MB of hard disk space, one 3.5-inch, 1.44 MB diskette drive, and a mouse or compatible pointing device.

10.3.2.3 Connectivity

IBM Internet Connection corporate dial services will allow switched communication speeds up to 28.8 kbps. The effective speed will depend on the type of modem and serial port the modem is connected to. A Hayes-compatible modem supporting 9,600 bps, or higher, and a telephone line are required. The following standards are currently supported:

- V.32 (9,600 bps)
- V.32bis (14,400 bps)
- V.34 (up to 28,800 bps)

The supported error control and data compression standards are:

- MNP Level 1-4 (error control)
- MNP Level 5 (data compression)
- V.42 (LAPM error control)
- V.42bis (data compression)

10.3.3 Leased Line Internet Connection Services

The Leased Line Internet Connection Services is part of the range of Internet services provided by the IBM Global Network. It offers a high-speed permanent and fully managed access link to the resources of the Internet. This service enables customers to conduct electronic commerce over the Internet by allowing them to provide information about their products and services and then actually sell them to customers if desired. Additional uses include:

- Interenterprise information exchange
- Electronic communication with business partners
- Corporate access to Internet databases

The IBM Global Network has more than 25,000 customer enterprises supporting more than 1.9 million users and access to networking services in 700 locations in nearly 100 countries. This network offers a spectrum of services designed to meet customers′ networking requirements for data, voice and video.

IGN provides leased line access to the Internet at speeds equivalent to corporate data networks. The services also expand the capabilities of IGN internetworking and multiprotocol solutions by allowing secure Internet access from their existing corporate networks.

10.3.3.1 Highlights

The Leased Line Internet Connection Services is the ideal solution for customers who want a permanent high-speed link to the Internet. They are available to customers using fully managed, dedicated communications facilities at speeds ranging from 56 kbps to 1.544 Mbps. This service is priced and operated to ensure the customer's business is able to leverage its information assets on the Internet with a complete, reliable and affordable service offering.

Capabilities include:

- Access for full TCP/IP connectivity to the Internet
- Managed dedicated leased line access to the Internet at high-speed data rates of 19.2, 56, 64, 128, 256, 512 kbps, 1.544 Mbps and 45 Mbps access on a special bid basis
- Assignment of IP address ranges for the customer network
- Assistance with registration of the customer private domain name with the responsible Naming Authority
- Fixed-price connections based on site connectivity requirements
- Internet Interconnect and IBM Global Network Firewall capabilities provide secure access from existing IGN internetworking and multiprotocol solutions to the internet

Leased Line Internet Connection Services offers customers managed full-time, high-speed access to the Internet via dedicated leased circuits. In the U.S., Advantis is offering two leased line access options to the Internet:

- Direct Leased Line Internet access provides an open two-way traffic between the customer's site and the Internet. No security is available.
- LAN Internetworking offers limited access to the Internet but it comes with firewall security for customers connected to the Advantis network environment through internetworking and multiprotocol solutions.

Customer Internet access requirements should be throughly reviewed to best choose the appropriate option. See your IBM Global Network local representative for additional information about requirements and availability of these offers in your country.

10.3.3.2 Features

IBM provides the planning, design, network components, installation, maintenance and operation required to attach customers' systems to IBM Global Network's Internet network.

The Leased Line Internet Connection Service includes:

- Backbone network, facilities and Network connectivity to the Internet through the IBM Global Network's Internet network.
- Customer premise router and backbone router(s).
- If required, an IBM 2210 Nways Multiprotocol Router for use as the customer site router (CSR), including an asynchronous modem for remote support/problem determination.
- Installation, maintenance and support of IBM-provided solution components.

- Data service units (DSUs)/customer service units (CSUs)
- LAN interface.
- Physical link (56 kbps-T1)
- If required, an IP address range for use in the customer's network will be assigned by IBM.
- Domain Name Services (DNS), where IGN will act as the external primary and/or secondary name server on behalf of a customer's network. IGN will negotiate with the Internet Network Information Center (NIC) or InterNIC to acquire network numbers as well as provide proper registration of IP addresses with the NIC on behalf of the customer and we will assist in connecting the customer's DNS to the global DNS infrastructure. This support is available immediately as part of the leased line Internet Connection capabilities.
- Network Management
 - 24-hour, seven-day-a-week network monitoring
 - Problem determination and management
 - Performance monitoring
 - Capacity planning and management of the IGN backbone network
 - Capacity monitoring of the CSR and circuit to the customer premise
 - Notification to the customer if an upgrade of the customer circuit is required
- Customer support
 - 24-hour, seven-day-a-week customer assistance

10.3.3.3 Physical Attachment Design

LAN Internetworking Version 1.1 offers firewall security protection via the IBM Global Network's product, TCPGATE2. It allows users with TCP/IP and/or SNA platforms to access limited Internet protocols. The supported features are Domain Name Server service, FTP, WWW browsing (via SOCKS gateway for TCP/IP users), Gopher, and Telnet. E-mail and Newsgroups support will be available in the future. Figure 190 on page 437 shows all network access paths to the IBM Global Network.

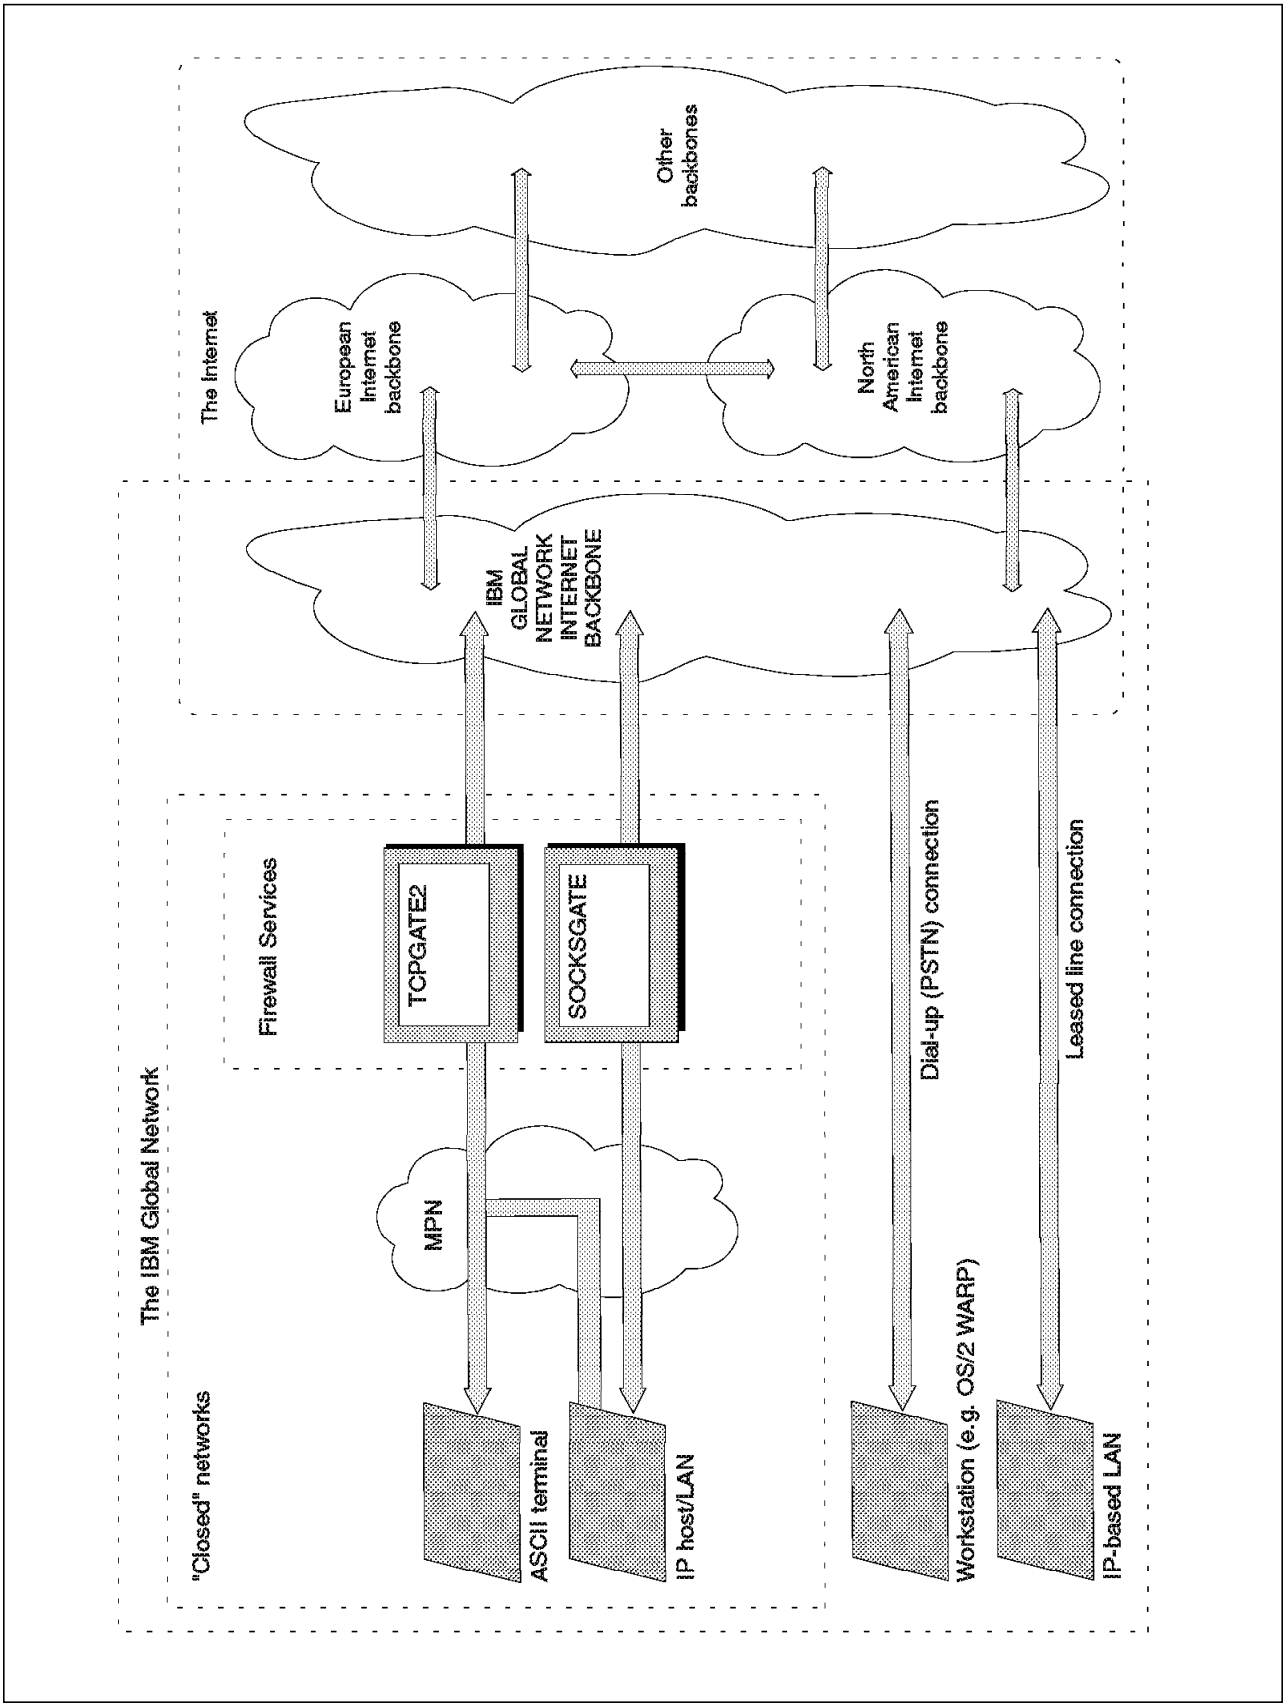


Figure 190. LAN Internetworking / Direct Leased Line via IBM Global Network

Direct leased line access provides a raw pipe of bandwidth between customer's site and the Internet. Users may choose to implement any Internet protocol on their own but no security is provided by IBM Global Network. Currently e-mail and Newsgroups support are not available via leased line offering, but users may employ their own. If you need additional information, refer to:

- *Leased Line Internet Connection Service - E/ME/A Attachment Guide* UH01-1003-00

The Leased Line Internet Connection Service (ICS) provides a permanent (non-switched) high-speed direct attachment to the IBM Global Network for customer's IP-based LANs (see Figure 191).

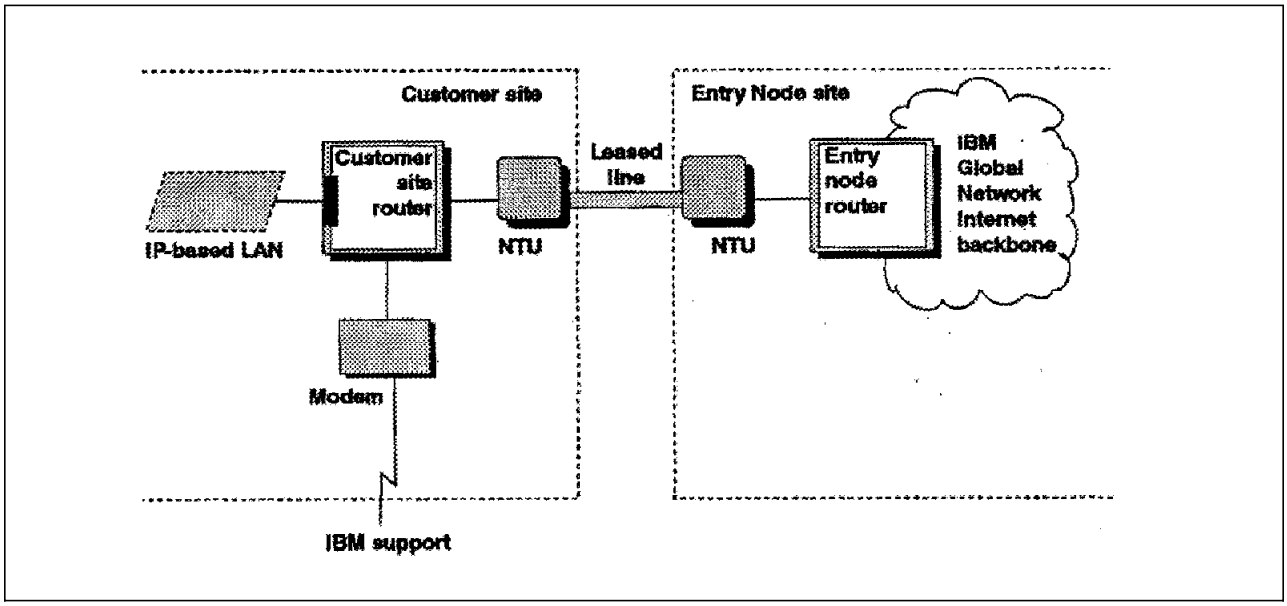


Figure 191. Direct Leased Line Internet Access Physical Attachment

The customer's LAN is attached, using a network interface card, to a customer site router (CSR). The CSR is then connected, via a leased line, to another router (the entry node router), which is directly connected to the IBM Global Network's Internet backbone (OpenNet). The CSR is also equipped with an analog dial-up port and a high-speed modem to allow IBM support personnel to access the CSR over the public switched telephone network (PSTN) to perform remote configuration, maintenance, and support.

10.3.3.4 Hardware and Software Requirements

IBM supplies and installs, if they are necessary, the following equipment at the customer site:

- A CSR with an appropriate network interface card to connect to the customer's LAN
- A PSTN modem and cables for use with the CSR's dial-up facility

Customers must provide:

- A TCP/IP enabled host and LAN, using the appropriate IP addresses.

- The appropriate cabling and connectors required to connect the customer's LAN to the network interface card on the CSR. The supported network types are:
 - Ethernet (10 Mbps)
 - Token-ring (4 Mbps and 16 Mbps)
- An analog PSTN circuit for use by the dial-up modem.

 Note: Customers planning to switch this circuit through a digital private automatic branch exchange (PABX), must ensure that the PABX is configured to provide an analog connection for the circuit. Customers with PABXs that do not support analog connections must ask the local PTT provider to supply a direct analog circuit for use by the dial-up modem.
- The leased line circuit from the customer site to the allocated IBM Global Network entry node. Where permitted by local legal and PTT regulations, IBM will order the appropriate leased line circuit on behalf of customers.
- The primary name server and its administration and support for names within the LAN. The primary name server should also be configured for inverse name address resolution.

 If required, IBM can supply the primary name server facilities for customers. However, a maximum of three network devices and two mail hosts only will be supported per customer.
- Security facilities, such as a firewall, to protect their network as required.

10.3.3.5 IP addresses

There are three classifications of IP addresses:

- Provider Aggregatable IP addresses (PA addresses)
- Provider Independent IP addresses (PI addresses)
- Private IP addresses (PR addresses)

PA addresses are globally unique addresses owned by an ISP (Internet Service Provider). When a customer terminates the contract with an ISP, any assigned PA addresses must be relinquished. The advantage to an ISP of using PA addresses for customer connections is that these addresses can be aggregated to a limited number of entries in the network routing tables. The advantages to customers is that the ISP can minimize the network routing tables, resulting in better performance.

PI addresses are also globally unique addresses, but are owned by customers. Customers can transfer these addresses from one ISP to another, provided that the new ISP is willing to support PI addresses. Unlike PA addresses, the routing of PI addresses through the Internet is not guaranteed; if the size of the network routing tables gets too large, ISPs may remove PI addresses from their tables. For this reason, the use of PI addresses is not recommended, and the use of PA addresses encouraged.

PR addresses are a range of addresses reserved by the Internet Assigned Numbers Authority (IANA) for use in private networks. That is, these addresses can be used in networks, provided that such networks do not have external connectivity. The disadvantage of using the addresses in this private address space is that when networks have to be merged, or when external

connectivity is required, then devices may need to be assigned new addresses; in some situations it may be possible to isolate the networks by using a firewall in between, but this is expensive in terms of the resources required.

10.3.3.6 IBM Global Network IP address policy

All customers attached to the IBM Global Network must use the correct classification of IP addresses depending on the type of connection. The following rules should be followed:

- If a firewall is installed between the customer's LAN and the IBM Global Network, then:
 - On the external side of the firewall:
 - PA addresses should be used.
 - PI addresses can be used, but at an additional charge.
 - PR addresses are prohibited.
 - On the internal side of the firewall:
 - PA addresses should be used.
 - PI addresses can be used, but at an additional charge.
 - PR addresses can be used, but customers should be aware of the disadvantages as detailed above.
- If a firewall is not installed between the customer's LAN and the IBM Global Network, then:
 - PA addresses should be used.
 - PI addresses can be used, but at an additional charge.
 - PR addresses are prohibited.

Note:

1. Customers who are using unregistered IP addresses and who do not plan to change to use either registered PA addresses or PI addresses have the following options:

- Install a firewall.
- Install an IP Address Translator.

In both cases, customers will still need to obtain either registered PA addresses or PI addresses for use on the Internet, but will not have to change the unregistered addresses currently used on their LAN.

2. Customers who already own PI addresses and who transfer these addresses for use with the ICS will be subject to a one-time charge due to the additional administrative effort required to support such addresses in the network routing tables.

For additional information about Leased Line Internet Connection Service, refer to:

- `http://www.ibm.com/globalnetwork/leasedbr.htm`
- *Leased Line Internet Connection Service - E/ME/A Attachment Guide*, UH01-1003-00

10.3.4 IGN′s Internet Backbone Design

IGN has backbone hubs in North America, Latin America, Europe, Africa, Asia, and Oceania.

- Asian Pacific OpenNet

 In Asia Pacific, eight backbone hubs have been implemented in addition to the eight Japanese cities that are connected. There are at least three more planned during 1996 (see Figure 192 and Table 32 on page 442).

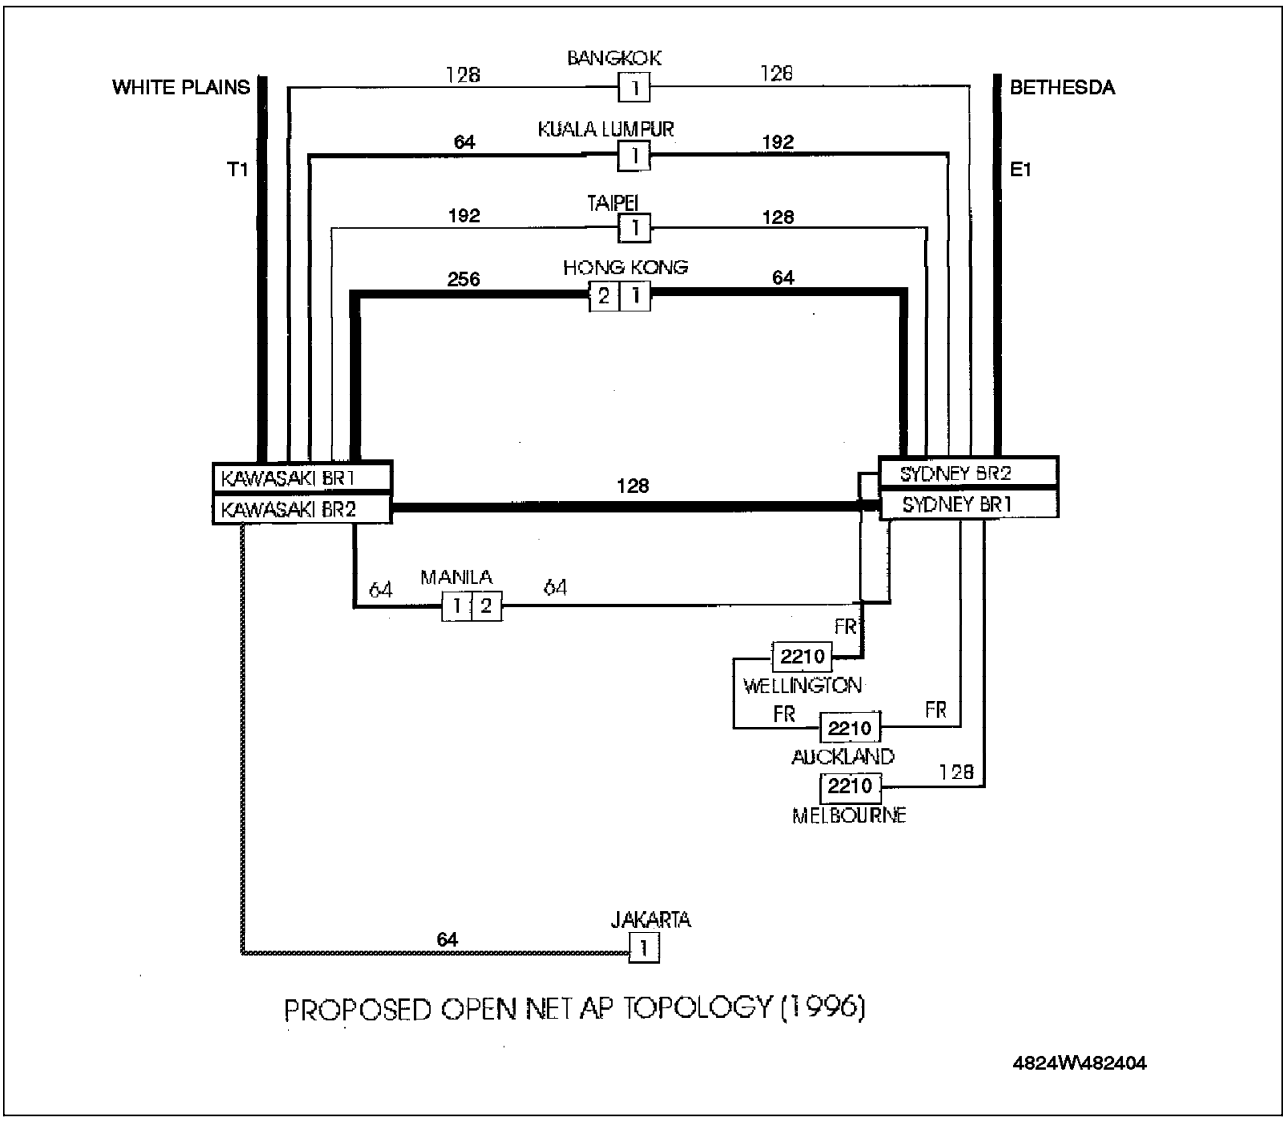


Figure 192. Asian Pacific OpenNet Node Sites

Table 32. AP OpenNet Node Sites (Excluding Japan Domestic-Only Nodes)

City	Node Type	Bandwidth to U.S.	Bandwidth: Sydney-Kawasaki	Bandwidth: to Sydney	Bandwidth: to Kawasaki
Kawasaki	Int'l Hub (2 X 6611)	1 X T1	1 X 128Kb		
Sydney	Int'l Hub (2 X 6611)	1 X E1	1 X 128Kb		
Hong Kong	country node (2 X 6611)			1 X 64Kb	1 X 256Kb
Melbourne	intra-country (1 X 2210)			1 X 128Kb	
Bangkok	country node (1 X 6611)			1 X 128Kb	1 X 128Kb
Jakarta	country node (1 X 6611)				1 X 64Kb
Taipei	country node (1 X 6611)			1 X 128Kb	1 X 192Kb
Kuala Lumpur	country node (1 X 6611)			1 X 192Kb	1 X 64Kb
Manila	country node (2 X 6611)			1 X 64Kb	1 X 64Kb
Wellington	country node (1 X 2210)			2 X FR•	
Auckland	country node (1 X 2210)			2 X FR•	

Note: • Frame Relay

Japan Domestic-Only OpenNet Nodes are:

- Tokyo
- Osaka
- Nagoya
- Fukuoka
- Hiroshima
- Sapporo
- Sendai
- Kanazawa

- EMEA OpenNet

 Throughout Europe, the Middle East, and Africa, the IBM Global Network has 29 major backbone hubs in 25 cities currently operational. IGN will deploy nine additional backbone hubs during 1996 (see Figure 193 on page 443, Table 33 on page 443, and Table 34 on page 444).

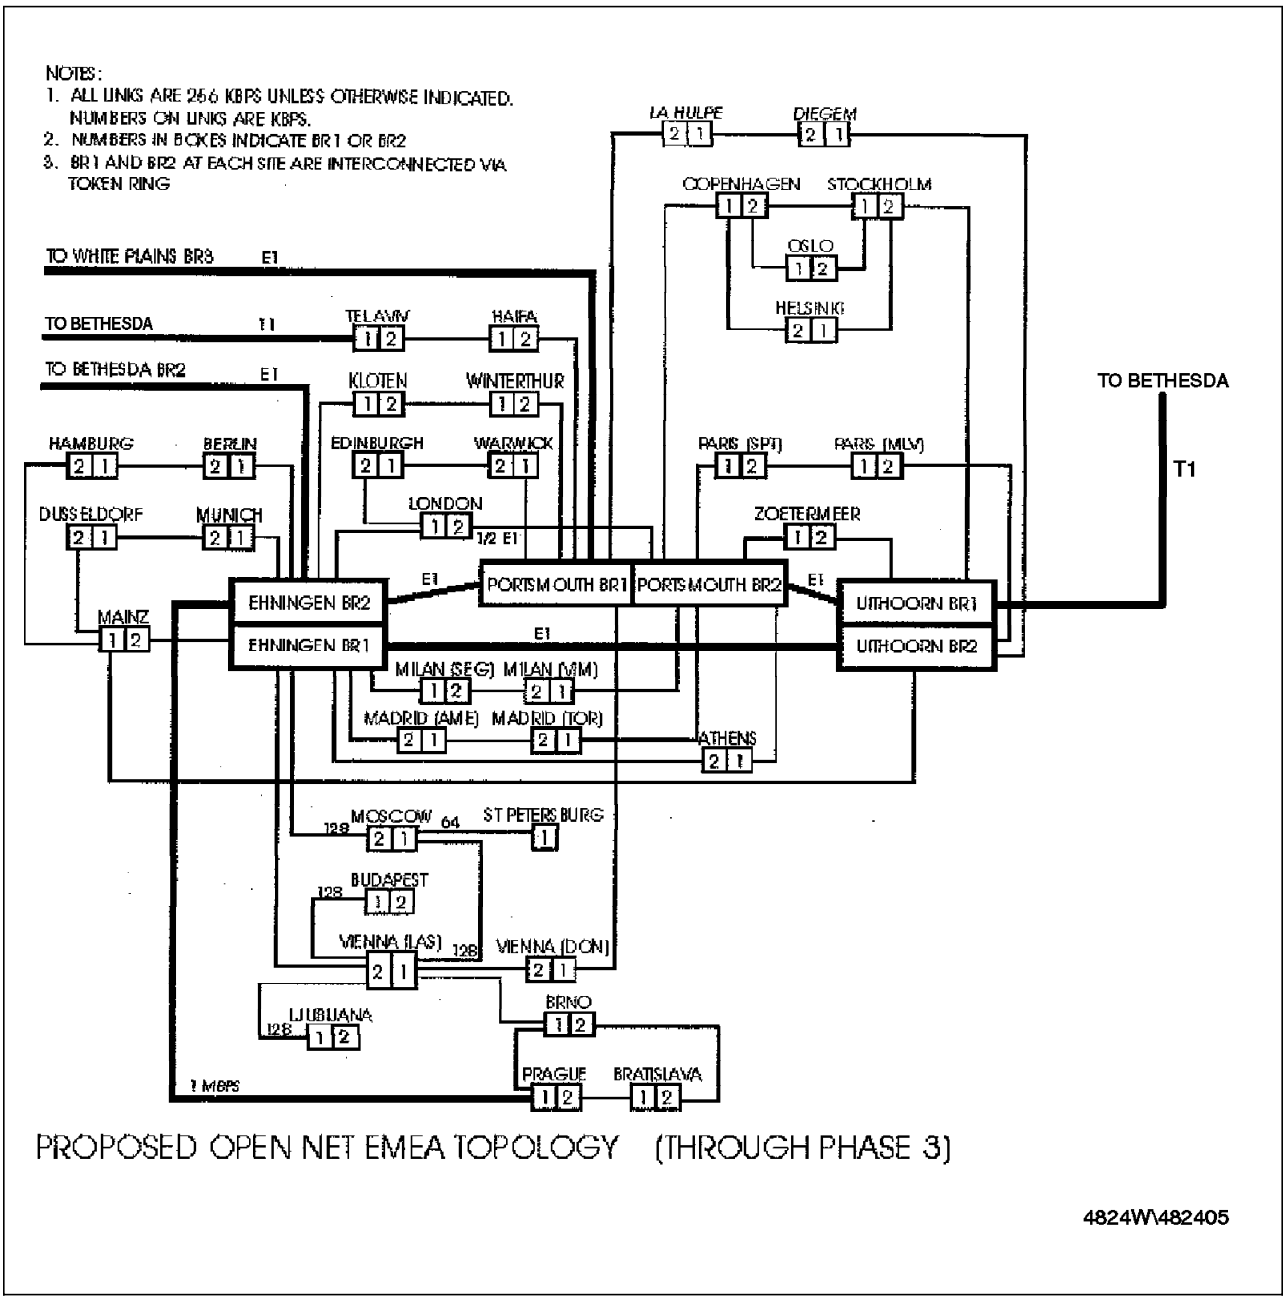


Figure 193. EMEA OpenNet Node Sites

Table 33. EMEA OpenNet Int'l Hubs

City	Node Type	Bandwidth to U.S.	Bandwidth: Ehningen-Portsmouth	Bandwidth: Ehningen-Uithoorn	Bandwidth: Portsmouth-Uithoorn
Ehningen	Int't Hub (2 X 6611)	1 X E1 to Bethesda	1 X E1	1 X E1	
Portsmouth	Int't Hub (2 X 6611)	1 X T1 to White Plains	1 X E1		1 X E1
Uithoorn	Int'l Hub (2 X 6611)	1 X T1 to Bethesda		1 X E1	1 X E1

Table 34 (Page 1 of 2). EMEA OpenNet Node Sites

City	Node Type	Bandwidth to	Link
Hamburg	country node (2 X 6611)	Berlin	1 X 256Kb
		Mainz	1 X 256Kb
Berlin	country node (2 X 6611)	EHNINGEN	1 X 256Kb
		Hamburg	1 X 256Kb
Dusseldorf	country node (2 X 6611)	Munich	1 X 256Kb
		Mainz	1 X 256Kb
Munich	country node (2 X 6611)	EHNINGEN	1 X 256Kb
		Dusseldorf	1 X 256Kb
Mainz	country node (2 X 6611)	EHNINGEN	1 X 256Kb
		UITHOORN	1 X 256Kb
London	country node (2 X 6611)	EHNINGEN	1 X 256Kb
		PORTSMOUTH	1/2 X E1
Edinburgh	country node (2 X 6611)	London	1 X 256Kb
		Warwick	1 X 256Kb
Warwick	country node (2 X 6611)	Edinburgh	1 X 256Kb
		PORTSMOUTH	1 X 256Kb
Kloten	country node (2 X 6611)	EHNINGEN	1 X 256Kb
		Winterthur	1 X 256Kb
Winterthur	country node (2 X 6611)	Kloten	1 X 256Kb
		PORTSMOUTH	1 X 256Kb
Tel Aviv	country node (2 X 6611)	WHITE PLAINS	1 X T1
		Haifa	1 X 256Kb
Haifa	country node (2 X 6611)	Tel Aviv	1 X 256Kb
		PORTSMOUTH	1 X 256Kb
La Hulpe	country node (2 X 6611)	PORTSMOUTH	1 X 256Kb
		Diegem	1 X 256Kb
Diegem	country node (2 X 6611)	La Hulpe	1 X 256Kb
		UITHOORN	1 X 256Kb
Copenhagen	country node (2 X 6611)	PORTSMOUTH	1 X 256Kb
		Stockholm	1 X 256Kb
Stockholm	country node (2 X 6611)	Copenhagen	1 X 256Kb
		UITHOORN	1 X 256Kb
Oslo	country node (2 X 6611)	Copenhagen	1 X 256Kb
		Stockholm	1 X 256Kb
Helsinki	country node (2 X 6611)	Copenhagen	1 X 256Kb
		Stockholm	1 X 256Kb
Paris (SPT)	country node (2 X 6611)	PORTSMOUTH	1 X 256Kb
		Paris (MLV)	1 X 256Kb
Paris (MLV)	country node (2 X 6611)	Paris (SPT)	1 X 256Kb
		UITHOORN	1 X 256Kb
Zoetermeer	country node (2 X 6611)	PORTSMOUTH	1 X 256Kb
		UITHOORN	1 X 256Kb
Milan (SEG)	country node (2 X 6611)	EHNINGEN	1 X 256Kb
		Milan (VIM)	1 X 256Kb

Table 34 (Page 2 of 2). EMEA OpenNet Node Sites

City	Node Type	Bandwidth to	Link
Milan (VIM)	country node (2 X 6611)	Milan (SEG)	1 X 256Kb
		PORTSMOUTH	1 X 256Kb
Madrid (AME)	country node (2 X 6611)	EHNINGEN	1 X 256Kb
		Madrid (TOR)	1 X 256Kb
Madrid (TOR)	country node (2 X 6611)	Madrid (AME)	1 X 256Kb
		PORTSMOUTH	1 X 256Kb
Athens	country node (2 X 6611)	EHNINGEN	1 X 256Kb
		PORTSMOUTH	1 X 256Kb
Moscow	country node (2 X 6611)	EHNINGEN	1 X 128Kb
		Vienna (LAS)	1 X 128Kb
Vienna (LAS)	country node (2 X 6611)	EHNINGEN	1 X 256Kb
		Vienna (DON)	1 X 256Kb
Vienna (DON)	country node (2 X 6611)	Vienna (LAS)	1 X 256Kb
		PORTSMOUTH	1 X 256Kb
Brno	country node (2 X 6611)	Prague	1 X 256Kb
		Vienna (LAS)	1 X 256Kb
Prague	country node (2 X 6611)	EHNINGEN	1 X 1Mb
		Brno	1 X 256Kb
Bratislava	country node (2 X 6611)	Prague	1 X 256Kb
		Brno	1 X 256Kb
St.Petersburg	country node (1 X 6611)	Moscow	1 X 64Kb
Budapest	country node (2 X 6611)	Vienna (LAS)	1 X 128Kb
Ljubijana	country node (2 X 6611)	Vienna (LAS)	1 X 128Kb

- Americas OpenNet

 The U.S. portion of the IBM Global Network contains 15 major backbone hubs. There are also three nodes in Canada and seven in Latin America, with additional expansion planned.

 For Latin America and Canada OpenNet node sites, see Figure 194 on page 446 and Table 35 on page 446.

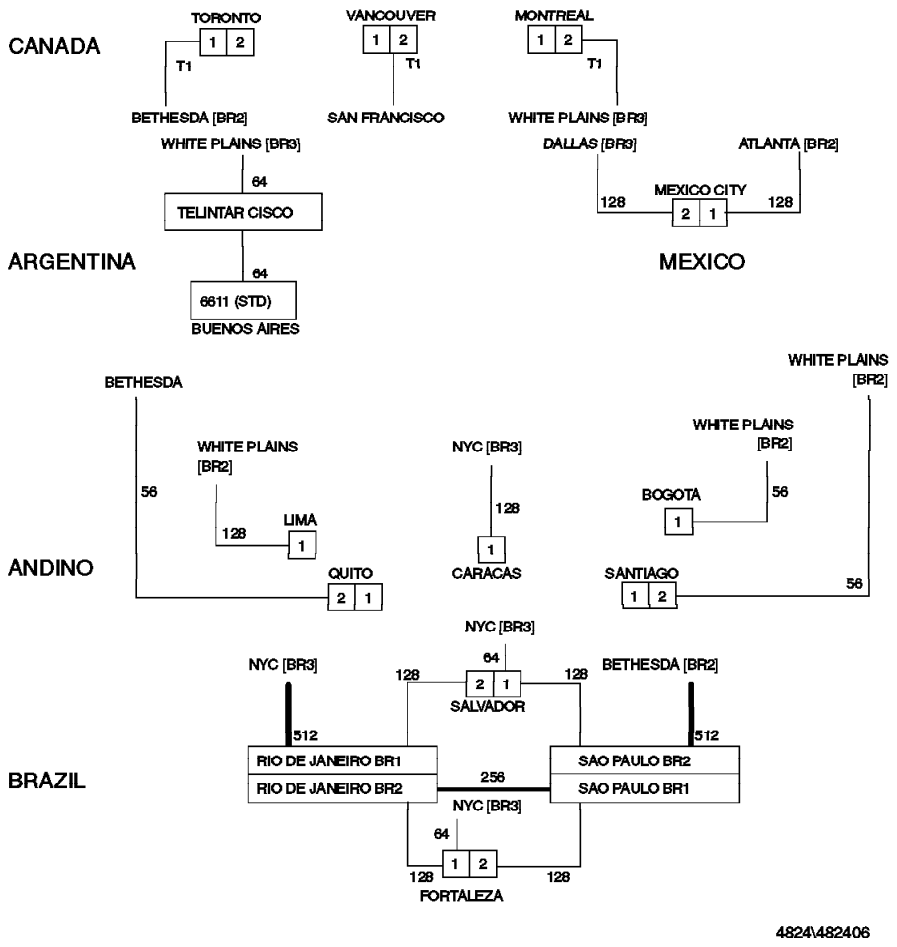


Figure 194. Latin America and Canada OpenNet Node Sites

Table 35 (Page 1 of 2). Latin America and Canada OpenNet Node Sites

City	Node Type	Bandwidth to U.S.		Bandwidth	To
		Bandwidth	To		
Montreal	Int'l Hub (2 X 6611)	1 X T1	White Plains		
Vancouver	Int'l Hub (2 X 6611)	1 X T1	San Francisco		
Toronto	Int'l Hub (2 X 6611)	1 X T1	Bethesda		
Sao Paulo	Int'l Hub (1 X 6611)	1 X 512Kb	Bethesda	1 X 256Kb	Rio de Janeiro
Rio de Janeiro	Int'l Hub (1 X 6611)	1 X 512Kb	New York City	1 X 256Kb	Sao Paulo
Salvador	Int'l Hub (1 X 6611)	1 X 64Kb	New York City	1 X 128Kb	Rio de Janeiro
Fortaleza	Int'l Hub (1 X 6611)	1 X 64Kb	New York City	1 X 128Kb	Rio de Janeiro
Quito	Int'l Hub (1 X 6611)	1 X 56Kb	Bethesda		

Table 35 (Page 2 of 2). Latin America and Canada OpenNet Node Sites					
City	**Node Type**	**Bandwidth to U.S.**		**Bandwidth**	**To**
		Bandwidth	**To**		
Santiago	Int'l Hub (2 X 6611)	1 X 56Kb	White Plains		
Bogota	Int'l Hub (2 X 6611)	1 X 56Kb	White Plains		
Lima	Int'l Hub (2 X 6611)	1 X 128Kb	White Plains		
Caracas	Int'l Hub (2 X 6611)	1 X 128Kb	New York City		
Buenos Aires	Int'l Hub (1 X 6611)	1 X 64Kb	White Plains via Telintar		
Mexico City	Int'l Hub (2 X 6611)	2 X 128Kb	Atlanta/Dallas		

For the U.S. OpenNet Topology, see Figure 195 and Table 36 on page 448.

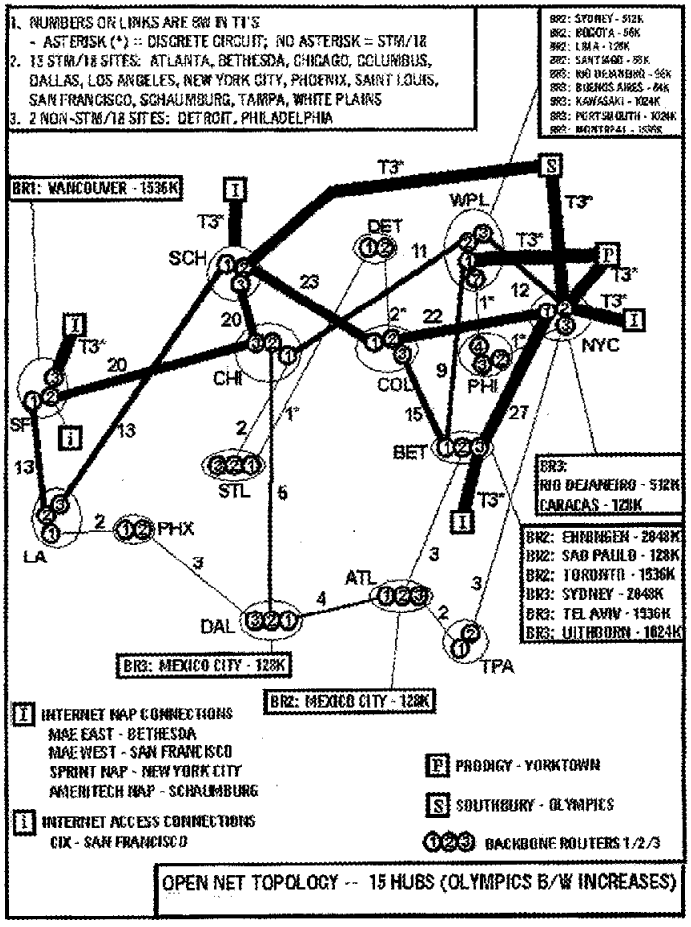


Figure 195. The United States OpenNet Topology

Table 36 (Page 1 of 2). OpenNet Topology

City/Hub	Link from/to	Bandwidth
Atlanta	Dallas	4 X T1 = 6Mb (STM/18)
	Bethesda	3 X T1 = 4.5Mb (STM/18)
	Tampa	2 X T1 = 3Mb (STM/18)
	Mexico City	128Kb
Bethesda	Columbus	15 X T1 = 22.5Mb (STM/18)
	White Plains	9 X T1 = 13.5Mb (STM/18)
	Atlanta	3 X T1 = 4.5Mb (STM/18)
	New York City	27 X T1 = 40.5Mb (STM/18)
	Mae East	1 X T3 = 45Mb
	Ehningen	2048Kb
	Sao Paulo	128Kb
	Toronto	1536Kb
	Sydney	2048Kb
	Tel Aviv	1536Kb
	Uithoorn	1024Kb
Chicago	White Plains	11 X T1 = 16.5Mb (STM/18)
	Saint Louis	2 X T1 = 2Mb (STM/18)
	Dallas	5 X T1 = 7.5Mb (STM/18)
	San Francisco	20 X T1 = 30Mb (STM/18)
	Schaumburg	20 X T1 = 30Mb (STM/18)
Columbus	Schaumburg	23 X T1 = 34.5Mb (STM/18)
	Detroit	2 X T1 = 2 X 1.5Mb (NON-STM/18)
	New York City	22 X T1 = 33Mb (STM/18)
	Bethesda	15 X T1 = 22.5Mb (STM/18)
Dallas	Atlanta	4 X T1 = 6Mb (STM/18)
	Chicago	5 X T1 = 7.5Mb (STM/18)
	Phoenix	3 X T1 = 4.5Mb (STM/18)
	Mexico City	128Kb
Detroit	Saint Louis	1 X T1 = 1 X 1.5Mb (NON-STM/18)
	Columbus	2 X T1 = 2 X 1.5Mb (NON-STM/18)
Los Angeles	Phoenix	2 X T1 = 3Mb (STM/18)
	San Francisco	13 X T1 = 19.5Mb (STM/18)
	Schaumburg	13 X T1 = 19.5Mb (STM/18)
New York City	Bethesda	27 X T1 = 40.5Mb (STM/18)
	Columbus	22 X T1 = 33Mb (STM/18)
	White Plains	12 X T1 = 18Mb (STM/18)
	Southbury	1 X T3 = 45Mb
	Prodigy/Yorktown	1 X T3 = 45Mb
	Sprint Nap	1 X T3 = 45 Mb
	Philadelphia	1 X T1 = 1 X 1.5Mb (NON-STM/18)
	Tampa	3 X T1 = 4.5Mb (STM/18)
	Rio de Janeiro	512Kb
	Caracas	128Kb

Table 36 (Page 2 of 2). OpenNet Topology

City/Hub	Link from/to	Bandwidth
Philadelphia	New York City	1 X T1 = 1 X 1.5Mb (NON-STM/18)
	White Plains	1 X T1 = 1 X 1.5Mb (NON-STM/18)
Phoenix	Los Angeles	2 X T1 = 3Mb (STM/18)
	Dallas	3 X T1 = 4.5Mb (STM/18)
Saint Louis	Detroit	1 X T1 = 1 X 1.5Mb (NON-STM/18)
	Chicago	2 X T1 = 3Mb (STM/18)
San Francisco	Los Angeles	13 X T1 = 19.5Mb (STM/18)
	IAC/CIX	
	Chicago	20 X T1 = 30Mb (STM/18)
	Mae West	1 X T3 = 45Mb
	Vancouver	1536Kb
Schaumburg	Los Angeles	13 X T1 = 19.5Mb (STM/18)
	Ameritech Nap	1 X T3 = 45Mb
	Southbury	1 X T3 = 45Mb
	Columbus	23 X T1 = 34.5Mb (STM/18)
	Chicago	20 X T1 = 30Mb (STM/18)
Tampa	Atlanta	2 X T1 = 3Mb (STM/18)
	New York City	3 X T1 = 4.5Mb (STM/18)
White Plains	Bethesda	22 X T1 = 33Mb (STM/18)
	Chicago	11 X T1 = 16.5Mb (STM/18)
	New York City	12 X T1 = 18Mb (STM/18)
	Philadelphia	1 X T1 = 1 X 1.5Mb (NON-STM/18)
	Prodigy/Yorktown	1 X T3 = 45Mb
	Sydney	512Kb
	Bogota	56Kb
	Lima	128Kb
	Santiago	56Kb
	Buenos Aires	64Kb
	Kawasaki	1024Kb
	Portsmouth	1024Kb
	Montreal	1536Kb

The IBM Global Network is a global supplier of Internet services, currently featuring more than 600 local Internet dial access points in nearly 50 countries worldwide. See Appendix E, “IBM Global Network Phone List” on page 595 for the IBM Global Network Phone List.

IGN also offers local dial numbers for online registration to access the Internet through IGN. See Appendix F, “IBM Global Network Registration Phone List” on page 611 for the IBM Global Network Registration Phone List.

IGN is always evaluating network access points (NAPs) to ensure high performance and reliability. IGN currently connects to five U.S. Interconnection points: Mae-East, Mae-West, Sprint NJ, and the Ameritech Chicago NAP, as well as to the Commercial Internet Exchange (CIX).

In Europe and the Middle East, IGN connects to the London Internet Exchange (LINX), the Belgian IP Interconnection Point (X-Router), the Israeli Internet Exchange (IIX), the Amsterdam Internet Exchange (AMS-IX), the MFS Frankfurt Exchange, the Vienna Exchange, the Stockholm Exchange (DGIX), and the French (GIX) Exchange in Paris. In Asia Pacific, IGN connects to the Hong Kong Internet Exchange (HKIX) and to the New Zealand Internet Exchange (NZIX). As more interconnection points emerge, IBM Global Network is positioned to connect to them.

Currently, IBM Global Network has redundant DS-3 access to the rest of the Internet.

For additional information, refer to URL:

`http://www.ibm.com/globalnetwork/inetbbon.htm`

Internet Operational Support

In the United States, a help desk is available 24 hours per day, 7 days per week via both an online problem management system as well as through a toll-free phone number. For help desk hours in other countries, check with your local support office. The network is monitored 24 hours a day and managed by network professionals. See Appendix G, “IBM Global Network Help Desk Phone List” on page 613 for the BM Global Network Help Desk Phone List.

For online problem management system, refer to URL:

`http://www.ibm.net/helpdesk.html`

Chapter 11. Content Services on the Internet

Internet content services can be described as the services performed to allow companies to respond quickly to the growing opportunity of doing business online using the Internet.

A company's presence on the Internet could be as simple as placing an electronic version of their executive brochure on a WWW server, or as complex as integrating customer service, ordering, marketing communications or other business processes with this electronic media.

Content services offer companies an opportunity for establishing a presence on the Internet using World Wide Web (WWW) technology. The customer provides and maintains the content, and the content services provides the space and the environment that is accessible to the users of the Internet.

The content services environment consists of multiple hosts (server workstations) attached via a LAN with direct, high-speed access to the Internet. Also, to become a content services provider you need to guarantee:

- Hardpware space and Software platform to host your customer's content.
- 24 hours a day, 7 days a week customer assistance to help identify and correct any problems that may occur.
- 24 hours a day, 7 days a week generally available service, except for scheduled maintenance.
- Domain Name Services (DNS), including registration of the customer's WWW domain name with the Internet Network Information Center.
- Activity reports to let the customer know how often network users access their content.

This chapter describes the content services concepts based on the IBM Global Network Content Services offering and guides the customer in how to create/implement a content hosting service in its own installation.

11.1 The Basic Internet Services

There are three basic Internet services: the World Wide Web, communication services, and information search and retrieval services. Depending on the service your customers will use, you have to set up and use specific servers, such as FTP, DNS, Mail, etc.

11.1.1 The World Wide Web

The most talked about and famous Internet service is the World Wide Web (WWW), which globally links documents together to form a *web* of information. Documents on the WWW can contain images, sound, clips, and even animation or video. The World Wide Web is the service that popularized the Internet.

The WWW links documents and transfers text, graphics, images, and voice information across the Internet using a special protocol called Hyper Text Transfer Protocol (HTTP). Documents and links are expressed in Hyper Text Markup Language (HTML). HTML also allows the author of a World Wide

© Copyright IBM Corp. 1996

Web to link to other documents Making home pages attractive, informative, and inviting is the key for a sucessful presence on the Internet.

11.1.2 Web Farms Concept

The Web farms concept is related to content hosting services, that is, the creation of customer's Web sites to provide key product or service information on servers connected to the Internet.

Web server farms also must be worldwide distributed and provide end-to-end management, systems operations and statistical reporting on Internet users who browse the customer's Web sites.

11.1.3 Communication Services

The Internet was originally designed for file transport between sites, so that researchers could share information and run their programs on other, faster computers. However, electronic mail (e-mail) and conferencing quickly became the most popular uses.

Today, the Internet is often used to exchange mail, and most electronic mail services (MCI Mail, America Online, Prodigy, CompuServe, etc.) can send and receive mail via the Internet, even if the Internet is not their native network. Mailing lists are an outgrowth of e-mail and contain the addresses of people with a common interest. There are thousands of mailing lists.

An alternative to mailing lists for people with common interest is the newsgroup. Think of a newsgroup as a bulletin board. You can either read the posted messages, add a message of your own, or comment on someone else's message. With a mailing list, the mail comes to you. With newsgroup, you have to go looking.

11.1.4 Information Search and Retrieval Services

Newsgroups and mailing lists handle notes and messages. What about files, such as programs, articles, pictures, and other larger collections of information? The Internet also provides services for these information types. The most basic way to find and retrieve information is via Telnet and FTP. With Telnet, you access a remote machine as a remote terminal user. If you can log on, you can do anything to the system within the capabilities the host machine provides.

FTP is more limited. FTP is designed specifically for file transfer. If the host machine has an FTP server, and you either have an account or the FTP supports anonymous access (using the special user name *anonymous*), you can log on and search the host's files for the information you want. With FTP, however, all you see is a collection of directories (or folders) and the files they contain. The first method developed to make FTP easier to use was a system called Archie. Archie, derived from the word archival, uses a central index of the files available on anonymous FTP sites around the Internet. Lists of file names are merged and can be searched for file names matching your target. Archie returns the locations of the names in the list that match your target. You then use FTP to retrieve them.

Searching with file names is cumbersome. Gopher severs were developed to simplify the process. Gopher provides menus for FTP, allows you to search

with keywords in addition to file names, and can help you link to other sites if the server you're linked to doesn't have what you need.

The WAIS (Wide Area Information Server) lets you ask for information in simple terms. Documents are indexed and keywords are placed into a WAIS database. This allows searches based on contents.

11.2 Content Services Concept

Based on the Web Farms concept, content services can be introduced as an information delivery service on the Internet.

The customers can host their information as Web pages, complete DBs or even complex applications using CGI programming (see Chapter 4, "Web Development" on page 175) and choose the way they want to make them available.

They can choose if they want Internet users to transfer their files through FTP services or just permit them to see the Web pages using browsers.

All of these steps depend on the content services of the provider's servers customization or depend on your own servers environment, if you want to install content services using your own installation.

11.3 Content Services through the IBM Global Network

IBM Content Services are provided by IGN - IBM Global Network, which provides support to customers wishing to access content services on the Internet, through IBM's worldwide network resources.

For further information about IGN, refer to the Chapter 10, "Connection Access Services" on page 419.

11.3.1 Highlights

IBM Content Services offer companies an opportunity to reach millions of new customers and prospects, market your products and services worldwide, and establish a presence on the World Wide Web without investing in new resources.

With content services through the IBM Global Network, you can distribute your company's information on the Internet easily, reliably and securely.

11.3.2 Enhanced Services

IBM content services offer enhanced services such as the following:

- Design and systems integration, including World Wide Web application and home page design
- Multimedia integration
- Data conversions and migration
- Content and server management
- Statistical information on how your Internet applications are used

- Around-the-clock network support, systems administration, backup and recovery, and security options

11.3.3 Versatility and Security

Through the IBM Global Network, all data safely resides on a server outside the internal network, so you can participate in the Internet marketplace and still maintain a secure environment. WWW uses the standard HTTP protocol to communicate and the standard HTML format to describe documents that reside on the servers.

WWW hypertext and information retrieval technologies pull together a powerful global information system.

11.3.4 Priced for Performance

With IBM Content Services, you pay only for what you use. This way your investment in the technology grows along with your potential customers′ acceptance of the medium.

The monthly charge is based on the amount of activity your server incurs for that month. Activity is defined as the number of requests satisfied within your server environment, that is, the number of *hits*.

11.3.5 Operating Environment

IBM supplies the appropriate hardware and software to host your WWW server. You are assigned an initial amount of megabytes of space for your information and provided with tools with which to define a staging area for testing and viewing of your home page before presenting it to the world.

IBM also provides:

- Customer assistance, 24 hours per day, 7 days per week, to help identify and correct any problems that may occur
- Generally available service 24 hours per day, 7 days per week
- Backup and recovery procedures to ensure the availability of your server

11.3.6 Connectivity to the Internet

Multiple high-speed links connect IBM′s Web farm (where your server resides) to IBM′s international high-speed Internet backbone. This technology used in IBM′s backbone is the same as that used in the NSFnet (National Science Foundation) backbone today. This backbone infrastructure is on a fast path to IBM′s ATM platform for the ultimate in performance and availability.

11.3.7 IBM Domain Name Services

The WWW domain name that you select will need to be registered with the Internet Network Information Center (InterNIC).

IBM does this for you and provides primary and secondary domain services, so that your users can easily find your home page.

11.3.8 Monthly Server Activity Report

IBM reports include:

- A summary section highlighting the number of requests that were made to access your content for the month.
- A detail section providing a daily and hourly view of the content activity.
- A summary of requests by domain name (for example, .com, .edu, and .org).

Further information about IBM Global Network Content Services is available via URL: `http://www.ibm.com/globalnetwork/contntbr.htm`

11.4 Creating a Content Hosting Service

In the following sections, you are going to see how to create a content service to make your customer′s information accessible on the Internet through the World Wide Web (WWW) and how to maintain your content utilizing the Web server′s environments.

You must be aware of everything, such as the content hosting description, network design sample, hardware and software platform′s considerations, a Web server SW installation, domain registration, etc.

11.4.1 Content Hosting Description

Content hosting means to host your customer′s information, DBs and applications using disk space on a server that is directly connected to the Internet.

This server hosts your customer′s company content, which can be accessed through a Uniform Resource Locator (URL) that they choose.

Depending on your customer′s demand, you are going to have one or more servers in the same network connected through a router to your networking provider or PTT (Post Telegraph Telephone - National Post and Telecommunication Authority).

If you are using a networking service provider, you must be connected via leased line if you intend to support applications on your servers.

For further information about leased line service, refer to the Chapter 10, “Connection Access Services” on page 419 and for information about routers, refer to Chapter 2, “Networking Hardware” on page 21.

11.4.2 Hardware Requirements

Web servers can be run from any hardware platform. In order to decide which hardware you must choose, many features must be analyzed. You must compare machines that offer the best technical features, such as memory size, HD size, speed, etc. Basically, these machines must be servers and fast ones.

The main hardware issue, therefore, is the amount of memory needed. Depending on what is going to be offered and made available, you may need more memory. If you are going to host just a few pages without graphics,

very little memory is required. On the other hand, if you intend to host and support pages with images, videos, sounds and large documents, you'll need a greater amount of memory.

You can use a PS/2, PC, RISC, AS/400 or an S/390 in your solution, but you must be aware of the number of your customers and the amount of data you need to keep or applications you need to run at the same time. This information and the size of your Web site and its network can determine which kind of machine is needed.

For further details about all possible HW solutions, refer to the Chapter 1, "Hardware Platforms" on page 1.

11.4.3 Software Requirements

Concerning software, you must have all the software necessary to run your content services.

Basically, you need to have:

- Operating system
- TCP/IP or TCP/IP stack
- Web browser
- Web server software
- Web server management software
- Web server report software

11.4.4 Connection Requirements

In networking you must be worried about your Web site link speed that must have at least a 56 kbps connection. This is the minimum acceptable for Web servers speed. Anything slower than this will immediately discourage users from accessing the site.

Faster connections (for example, a T-1 line) are also more expensive; to find the balance between the cost of a connection and your company's budget.

Another solution to consider is the service provider. In this case you only need to pay for a leased line circuit to your service provider in order to connect your Web site on the Internet. To choose your connection service provider, you must consider three important factors:

- Cost of the services
- Accessability
- Reliability

For further details about connection services provider, refer to the Chapter 10, "Connection Access Services" on page 419.

11.4.5 Network Solution Design Sample

Based on all of these requirements, Figure 196 shows a sample of a basic content services network solution design that you can consider when building your own service.

We are representing the Web server hardwares, which depend on your solution design (that is, the services you want to offer).

We also show a Web site workstation dedicated to the administration service and a router to connect your LAN to the Internet directly or through a connection service provider.

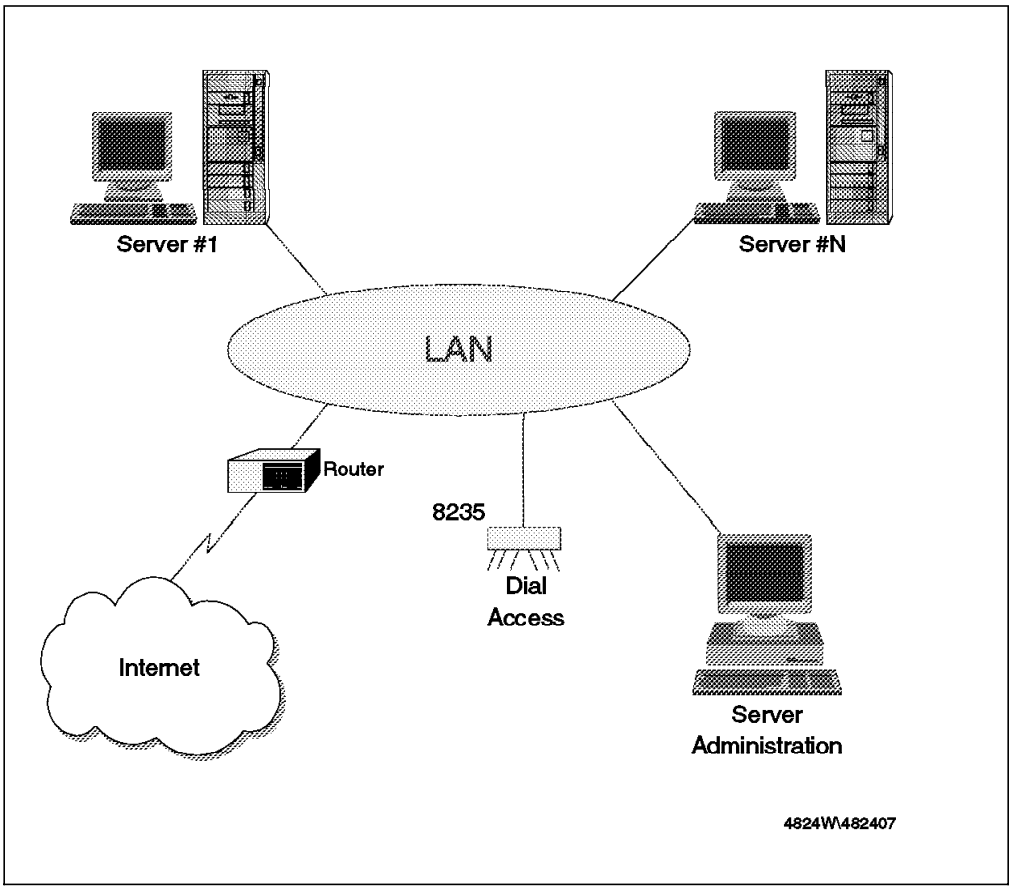


Figure 196. Internet Content Services Network Environment

11.4.6 IP Addressing

The Internet is comprised of both physical wires and software connections. When you try to imagine what the Internet is and how it operates, it is natural to think of a chaotic unmanaged network. How does a single request know where to go? This is where an Internet address or IP address is used.

The IP address is based on a hexadecimal numbering system. The clever part of the IP address is that the numbers are chosen to make the network and routing more efficient. Specifically, an IP address encodes the identification of the network to which an end user is attached within the IP address specified at the IP network layer.

Every interface on the Internet must have a unique IP address. This chapter will not go into the complexities involved in designing an IP network. However, to be able to understand the domain concept we are introducing, some of the basics of IP addressing need to be understood.

Each host attached to the Internet has an assigned unique 32-bit universal identifier, or IP address. Conceptually, each IP address is made up of a pair of numbers: the network ID (net ID) and host ID (host ID). In practice, this pairing can take one of three classes, as follows:

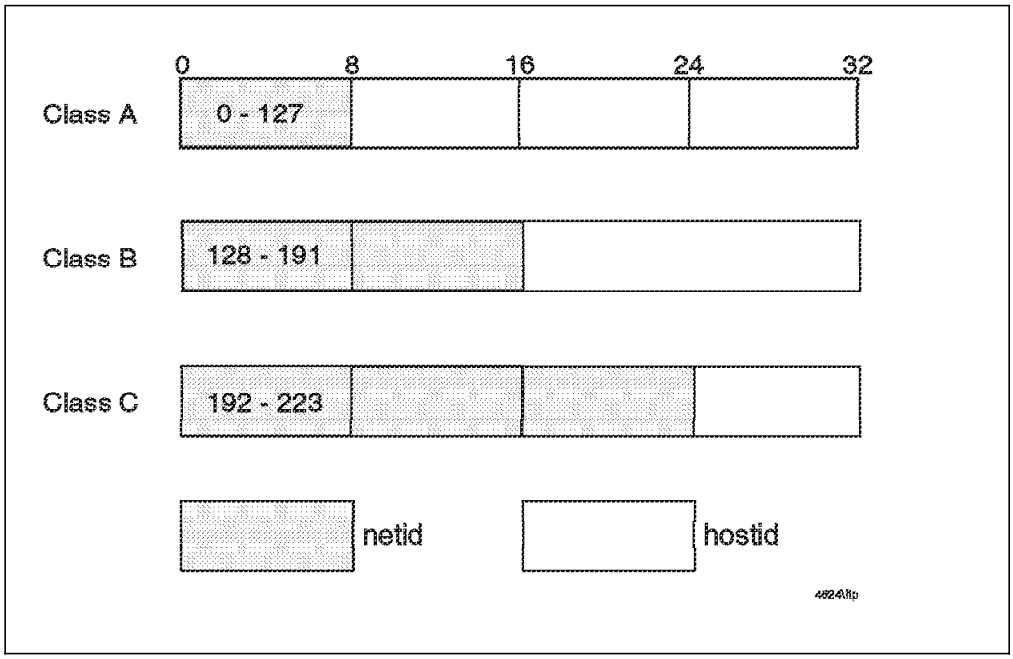


Figure 197. IP Address Classes

Each network class will allow different possible network and host combinations, as shown in Table 37.

Table 37. Class versus Network and Hosts

Class	Number of Networks	Number of Hosts
A	Less than 256	Greater than 65536
B	256 to 65536	
C	Greater than 65536	Less than 256

For the ease of communicating, IP addresses are written as four-decimal integers separated by decimal points, where each integer is given the value of one octet of the IP address. Thus a 32-bit address is written as xx.xx.xx.xx. For example, the binary network address:

8 16 24 32

10000000 00001010 00000010 00011110

is written:

128 10 2 30

or:

128.10.2.30

Since every host on the Internet must have a unique IP address, there must be some central authority for allocating these addresses for networks and hosts. This authority is the Internet Network Information Center (InterNIC).

InterNIC is responsible for network and domain registration. End users do not get their IP address from InterNIC. InterNIC normally assigns a range of IP addresses to service providers. To get an IP address, you must approach your service provider, who, depending on your connection type, will assign you an IP number from a range of IP addresses that they have been allotted.

If you do not want to connect through a service provider and intend to connect to the Internet directly, you must apply to InterNIC for a domain address and an IP network ID. To apply directly to the InterNIC, you must be either a service provider or a very large global corporation. The assignment of host IDs is then up to the system administrator on your site.

InterNIC does not readily provide a direct service and will, in almost every case, redirect queries through to a service provider. Two classes of service providers exist. Some service providers operate at a regional level and are responsible for a wider range of top-level IP addresses. This is covered in more detail in RFC 1466.

Further information about InterNIC registration can be found at the URL: http://www.internic.net or via e-mail at info@internic.net.

11.4.7 Domain Name Systems

In the TCP/IP world, the Domain Name System (DNS) is a distributed database system that provides the mapping between IP addresses and host names. We use the term distributed because no single site on the Internet knows all the information. Each site maintains it own database and runs a database or name server that other systems accross the Internet can query. The DNS provides a protocol that allows clients and servers to communicate with each other.

In 1992, the Internet Architecture Board (IAB) wrote to the Defense Information Systems Agency (DISA) regarding the phasing out of the old host name to address tables and the wider adoption of the Domain Name System (DNS). This correspondence marked the end of a system that had first been adopted in the early 1980s by the Department of Defense (DoD) and the DDN Network Information Center (NIC).

11.4.7.1 Name Systems

The IP protocol requires its 32-bit IP network address for each host. Token-ring and Ethernet technologies require unique hardware or MAC (Media Access Control) addresses for the interfaces onto the cable. Now, as users of these protocols and physical technologies, we need to use the addresses to communicate. But people are not very good at remembering large numbers of 32-bit IP addresses or 48-bit MAC addresses. We use telephone numbers all the time, but we don't try and remember each and every one of them. Instead we use a directory. This is a list that maps the name of the person we want to contact to their telephone number. This is exactly the problem that faced the growing numbers of Internet users. How

do you remember the individual addresses of each of the hosts on the Internet?

11.4.8 The Flat Name Space

The initial answer was a simple one: the Internet Host Table. Specified in RFC 810 - DoD Internet Host Table Specification, the Internet Host Table was a flat file that was maintained by the NIC. Each host registered its symbolic name and IP address with the NIC, and the NIC updated its HOSTS.TXT table. Users would then obtain a copy of the file via FTP from the NIC host.

RFC 810 - DoD Internet Host Table Specification laid down a specification for the structure of the host names as they would be used in the table, defining each as an ASCII text string with six fields separated by colons. Each entry is then defined as either a NETWORK, GATEWAY or HOST entry, with additional comments relating to the type of hardware, operating system and protocols that this particular host employed. An example of the host table format would appear as follows:

```
NET : 10.0.0.0 : ARPANET
NET : 128.10.0.0 : PURDUE-CS-NET :
GATEWAY : 10.0.0.77, 18.10.0.4 : MIT-GW.ARPA,MIT-GATEWAY : PDP-11 :
          MOS : IP/GW,EGP :
HOST : 26.0.0.73, 10.0.0.51 : SRI-NIC.ARPA,SRI-NIC,NIC : DEC-2060 :
       TOPS20 : TCP/TELNET,TCP/SMTP,TCP/TIME,TCP/FTP,TCP/ECHO,ICMP :
HOST : 10.2.0.11 : SU-TAC.ARPA,SU-TAC : C/30 : TAC : TCP :
```

This *flat name space* approach appeared to resolve the initial problem. So what went wrong?

11.4.8.1 The Name Space Explosion

In 1987 it was recognized that the continued growth in the Internet was causing problems to the name/address translation services. The bandwidth required to transfer the HOSTS.TXT file to all the hosts on the Internet was proportional to the number of hosts on the Internet and was increasing rapidly. The types of hosts out on the network were also changing. Local networks were emerging with organizations administering their own addresses and names. Local changes to this administration could be made at will, but there was a delay before the NIC could update its HOSTS.TXT file and ship it out to the rest of the Internet. The applications running on these hosts were becoming more and more sophisticated and there was an increasing need for a general purpose name service with an element of local structure to give organizations more flexibility and control. The answer was the Domain Name System (DNS).

11.4.9 The Domain Name System

A variety of proposals emerged to counter the problems of the flat name space, but each of them suggested a hierarchical name space using a distributed database. The hierarchical approach would allow for the delegation of authority and provide organizations with the level of control they required. The distributed database would ease the problems of size of the database and the frequency of its updates. The resulting scheme, DNS, has three major components:

- The domain name space and resource records specify the hierarchical name space and the data associated with the resources held within it.

Queries to the name space extract specific types of information from the records for the node in question.

- Name servers are server programs that hold information about the name space structure and the individual sets of data associated with the resources within it.
- Resolvers are programs that extract information from the name servers in response to client requests.

We begin our discussion of DNS with a look at each of these elements in turn.

11.4.9.1 The Domain Name Space

The domain name space is essentially a distributed database containing information about the hosts and gateways in the Internet. Not only does it provide a mapping of the IP address to a symbolic name for the host, but it also offers information on the resources available on that host, such as its hardware, operating system and the protocols and services in use.

The name space is built as a hierarchical tree structure with a root at the top. This root is unnamed and is often represented by a single period (.). The tree has branches, each emanating from an intersection point called a *node*. Each node corresponds to a resource (a host or gateway).

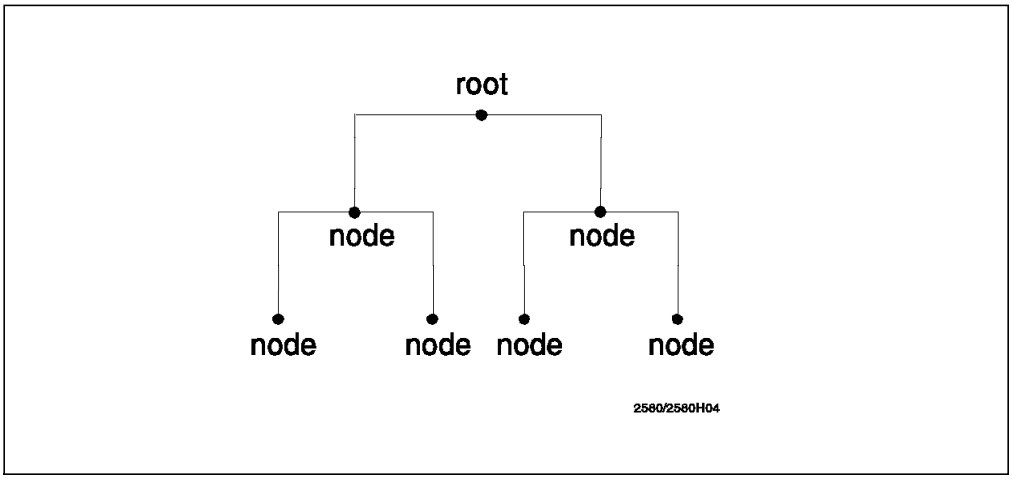


Figure 198. The Tree Structure of the Domain Name Space

We have called this structure the domain name space, but what exactly is a domain? A domain is identified by a domain name. It consists of the part of the name space structure which is at or below the domain name. Thus, a domain starts at a named node and encompasses all those nodes that emanate below it. Let us look at an example:

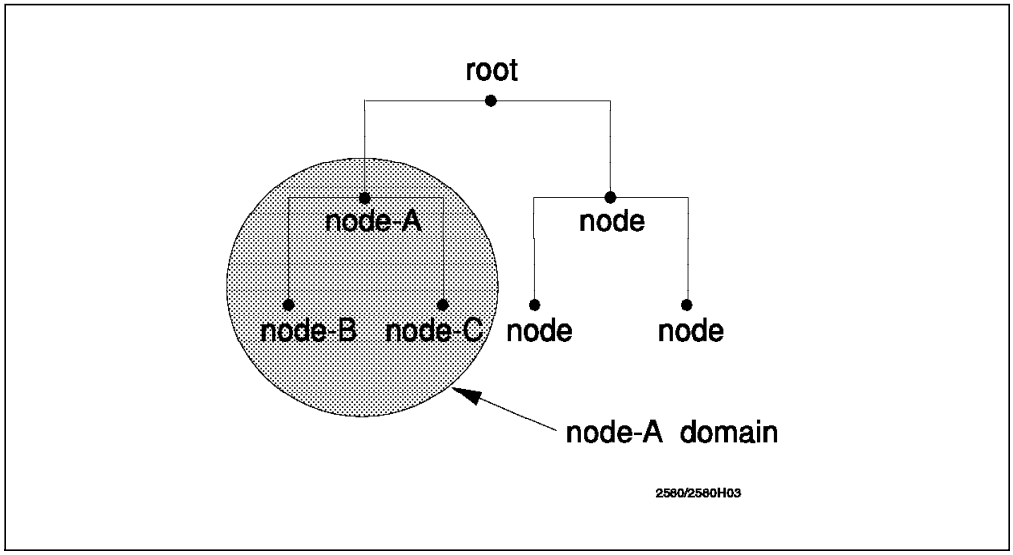


Figure 199. The DNS Domain

This shows a domain node-A, that begins at node-A. It contains node-A, node-B and node-C. This scheme may be taken a step further to show that as we progress out from the root, we will create *subdomains*. The next diagram illustrates this.

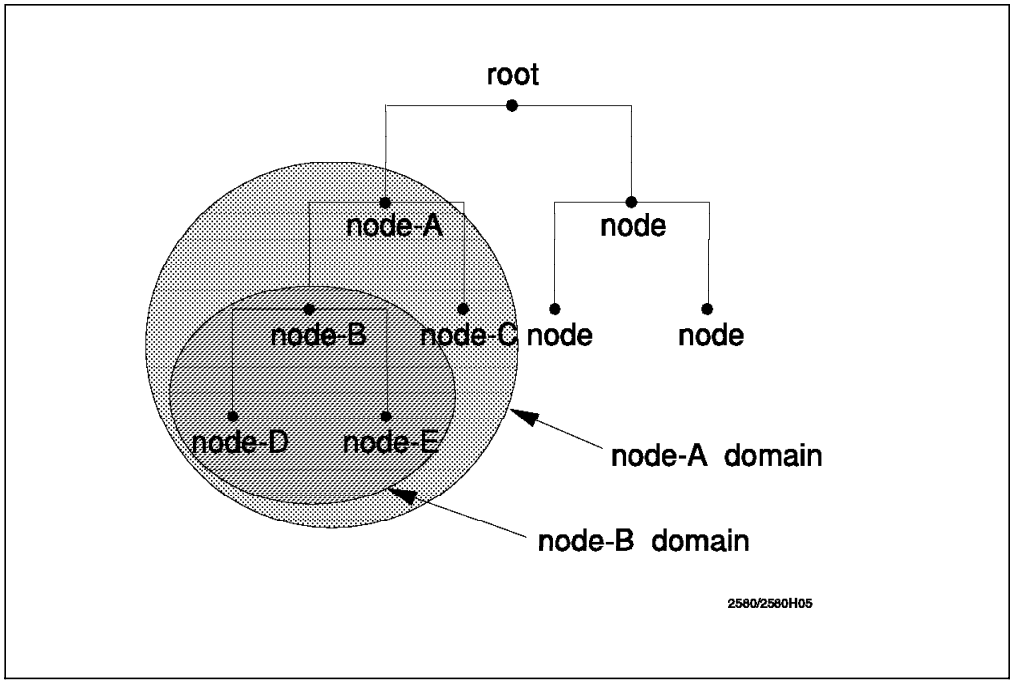


Figure 200. DNS Subdomain

A new domain, node-B, contains node-B, node-D and node-E. The original domain, node-A, now encompasses not only node-A, node-B and node-C but also the subdomain created by node-B.

Domain Names

Each node in the tree is labeled with a name of up to 63 characters in length. This label must start with a letter, end with a letter or digit and contain only letters, digits or hyphens (-). For example:

```
SRI-NIC    (the Network Information Centre at SRI International)
```

Currently, domain names are not case sensitive. A node may have a label AAA which could be referred to as either AAA or aaa . However, it is strongly recommended that you preserve the case of any names you use. Some operating systems, UNIX for example, are case sensitive, and future developments of the DNS may possibly implement case-sensitive services.

The name does not have to be unique in itself; some names appear many times in the name space. However, to ensure that each node in the tree can be uniquely identified, it is stipulated that sibling nodes (that is, those nodes with the same parent node) must not use the same name. This limitation applies only to the child nodes, and the name may appear in a node with a different parent.

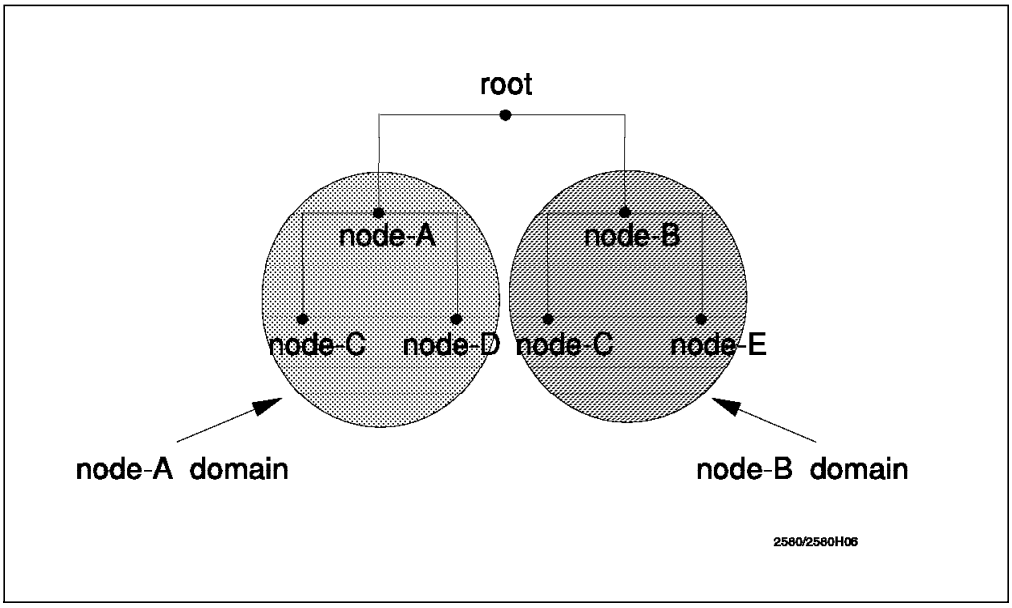


Figure 201. Domain Names

Figure 201 illustrates how a name may appear more than once within the tree. The name node-C appears twice in the tree (once as part of the domain node-A and again as part of the domain node-B). Node-A and node-B are siblings (have the same parent node, which is *root*) and so their names must be unique. Node-C and node-D in the node-A domain are also siblings and must again be named uniquely. However, node-C in the node-B domain has a different parent node to node-C in the node-A domain. To maintain the unique identity of each node, it is therefore apparent that we must use the identity of its parent node whenever we reference a node outside of its own domain. This scheme *qualifies* the name and provides what is known as a *fully qualified domain name (FQDN)*.

Fully Qualified Domain Name (FQDN)

The use of an unqualified name within a domain is the efficient way that names are used in preference to addresses and is perfectly valid. For example, referring to USER1 is much easier for remembering than using the 32-bit IP address 172.16.3.14 . However, the IP address is unique within the Internet while the name node-C (as we have shown previously) may not be. The answer is the FQDN. To create the FQDN of a node we must use the sequence of names on the path from the node back to the root with periods separating the names. These names are read from left to right, with the most specific name (the lowest and farthest from the root) being on the left. Thus, we see that the two hosts in our previous example now have completely unique FQDNs:

```
node-C.node-A.root      and      node-C.node-B.root
```

In practice, the name of the root domain is never shown; it has null length and is usually represented by a period (.). When the root appears in a domain name, the name is said to be *absolute*. For example:

```
node-C.node-A.    (The root is represented by the trailing period.)
```

This makes the FQDN totally unambiguous within the name space. However, domain names are usually written relative to a higher level domain rather than to the root itself. In the previous example, this would mean leaving off the trailing period and referring to node-C relative to the node-A domain. For example:

```
node-C.node-A
```

When you configure a TCP/IP host, you are requested to enter the host name of the host and the domain origin to which this host belongs. In the previous example, if we configured a host in the node-C.node-A domain, we would enter the host name as, for example, host-X and the domain origin as node-C.node-A. Whenever a nonqualified name is entered at this host, the resolver will append the current domain origin to the name, resulting in an FQDN belonging to the same domain as our own host, which enables us to refer to hosts that belong to the same domain as this host, by just entering the unqualified host name. If we enter host-Y, the resolver will append the domain origin building the fully qualified name host-Y.node-C.node-A before trying to resolve the name to an IP address. If we want to refer to hosts outside our own domain, we will enter the fully qualifed name as, for example, host-Z.node-E.node.A.

Top-Level Domain (TLD)

There is seemingly no restriction on the names that you can create for each node, other than that of length and uniqueness among siblings. However, the NIC decided to provide some sort of order within the name space to ease the burden of administration. Below the root are a number of *top-level domains* or *(TLDs)*. These TLDs consist of seven generic domains established originally in the USA to identify the types of organization represented by the particular branch of the tree. These can be seen in Figure 202.

United States Only Generic Domains

gov - Government institutions - now limited to US Federal agencies
mil - US Military groups only

Worldwide Generic Domains

edu - Educational institutions
com - Commercial organizations
net - Network providers (such as NSFNET)
int - International organizations (such as NATO)
org - Other organizations that do not fit anywhere else

Figure 202. The Generic Top-Level Domains

The generic TLDs first outlined for the Domain Name System were augmented by the 2-character international country codes as detailed in the ISO 3166 standard. Known as *country* or *geographical* domains, these TLDs often have subdomains that map to the original US generic top-level domains such as .com or .edu.

11.4.9.2 Domain Name System Resource Records

We have looked at the structure of the domain name space and discussed nodes and resources. Each node is identified by a domain name and has a set of resource information composed of *resource records (RRs)*. The original concept of the name system was to provide a mapping of names to addresses, but it has proved far more useful than just that. The resource records contain information about the node: the machine type it is running on, the operating system and services it runs, and, more importantly, information about the mail exchange within the domain.

The format of a resource record and a description of each term is shown below:

```
name ttl class type rdata
```

name This is an owner name, that is, the domain name of the node to which this record pertains (maximum length is 255 characters).

ttl This is the *time-to-live*. This is a 32-bit unsigned value in seconds that this record will be valid in a name server cache. A zero value means the record will not be cached but will be used only for the query in progress. This is always the case with SOA records.

class This is the class of the protocol family. The following values are defined:

Class	Value	Meaning
-	0	Reserved
IN	1	The Internet
CS	2	The CSNET class (now obsolete)
CH	3	The CHAOS class
HS	4	The Hesiod class

type This is the type of resource defined by this record. The following values are defined:

Type	Value	Meaning
A	1	A host address
NS	2	The authoritative name server for this domain
CNAME	5	The primary (canonical) name for an alias
SOA	6	Marks the start of a zone of authority in the domain name space
WKS	11	Describes the *well-known services* that are supported by a particular protocol on this node, TCP(FTP), for example
PTR	12	A pointer to an address in the domain name space; used for address to name resolution
HINFO	13	Information about the hardware and operating system of this node
MX	15	Identifies the domain name of a host which will act as a mailbox for this domain
TXT	16	Text strings

rdata This is the data associated with each record The value depends on the type of value defined, with most types having several elements:

Type	Rdata value
A	A 32-bit IP address (for the IN class)
NS	A domain name
CNAME	A domain name
SOA	The domain name of the primary name server for this zone. A domain name specifying the mailbox of the person responsible for this zone An unsigned 32-bit serial number for the data in the zone, usually in the format (yyyymmdd) A 32-bit time interval before the zone is refreshed (seconds) A 32-bit time interval before retrying a refresh (seconds) A 32-bit time interval before data expires (seconds) An unsigned 32-bit minimum TTL for any RR in this zone
WKS	A 32-bit IP address An 8-bit IP protocol number A variable length bit map (multiples of 8 bits long) with each bit corresponding to the port of the particular service
PTR	A domain name
HINFO	A character string for CPU type (see list in RFC 1700) A character string for operating system type (see list in RFC 1700)
MX	A 16-bit integer specifying the preference given to this RR over others at the same owner (lower values are preferred) A domain name
TXT	One or more character strings

An example of these resource records is given in the following section.

Sample DNS Master File

The sample network we created has a number of nodes, as seen in Figure 203.

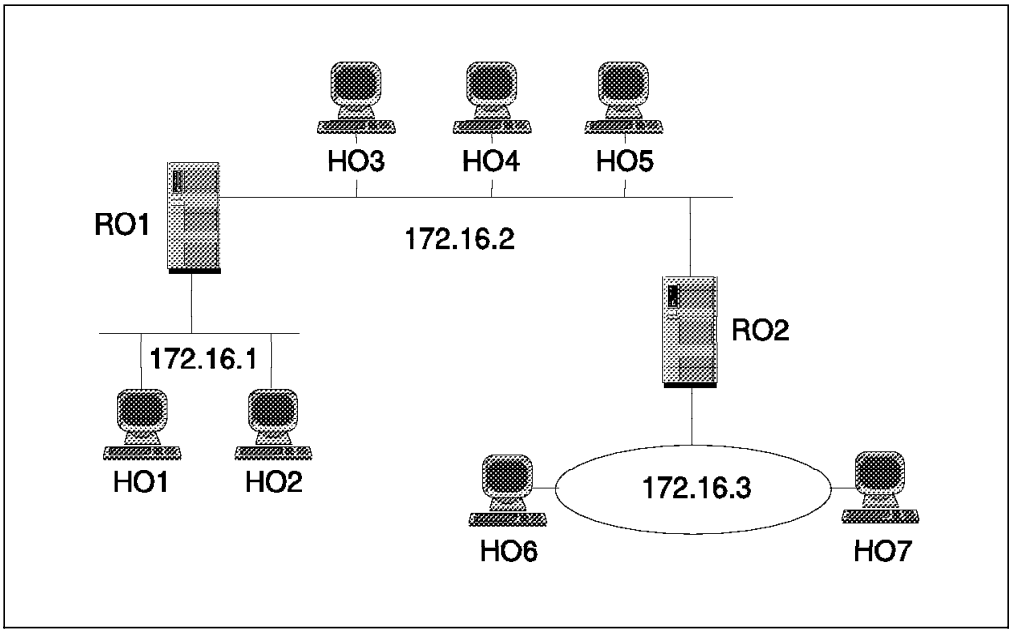


Figure 203. A Sample Network

This sample network contains three physical networks connected by two routers. We have subnetted our IP network number to provide connectivity through the routers. Host HO5 has been assigned the task of name server. We have created a single domain to cover all the hosts in all three networks, with a domain name of sample.net. At this point you will notice that although we have three physical networks, we only need a single domain. The domain is a logical concept and bears no relationship to the physical networks it covers. However, it would of course be possible to create three subdomains (one for each of the subnets) if that made it more efficient to administer. We deal with this aspect of administration later.

These resource records are stored in text format in a file called the *master file*. This is used as input to the actual database that holds the information on the name server. The format of the master file is a sequence of line-oriented entries, with parentheses as continuation characters. Comments are denoted by lines which start with a semicolon (;).

There are two control entries defined: $origin and $include. $origin is used and explained in the following example. $include (filename) is not seen in this example, but allows you to insert the named file into the current master file.

```
;
$origin sample.net.  1
;
@   2        IN SOA HO5.sample.net. JIM.HO5.sample.net.  (
                       19950517          ;serial number for data
                       10800             ;secondary refreshes every 3 hour
                       3600              ;secondary retries every 1 hour
                       604800            ;data expire after 1 week
                       86400)            ;minimum TTL for data is 1 day
;
@       99999 IN NS  HO5.sample.net.  3
;
HO5     99999 IN A   172.16.2.3   4
        99999 IN WKS 172.16.2.3 TCP (SMTP   5
                                     FTP
                                     TELNET
                                     NAMESRV)
;
HO4     99999 IN A     172.16.2.2
              IN HINFO IBM-PS/2/OS/2   6
HO3     99999 IN A     172.16.2.1
              IN HINFO IBM-PS/1/PCDOS
HO1     99999 IN A     172.16.1.1
        99999 IN MX  0 HO1   7
        99999 IN MX  5 HO2   7
HO2     99999 IN A     172.16.1.2
        99999 IN MX  0 HO2
        99999 IN MX  5 HO1
HO6     99999 IN A     172.16.3.1
HO7     99999 IN A     172.16.3.2
;
;RO1 and RO2 are routers and each have 2 different IP addresses
;
RO1     99999 IN A     172.16.1.3
        99999 IN A     172.16.2.4
        99999 IN TXT   IBM 6611 located on 1st floor   8
RO2     99999 IN A     172.16.2.5
        99999 IN A     172.16.3.3
        99999 IN TXT   IBM 6611 located on 2nd floor
;
; Aliases
;
host2   99999 IN CNAME HO2   9
host7   99999 IN CNAME HO7
;
```

Figure 204. DNS Master File on the Name Server HO4

Notes:

1 The $origin statement identifies the origin of the zone (sample.net. in our case). This name will be appended to all the resource names in the master file that do not end with a period. For example, HO4 will become a fully qualified domain name of HO4.sample.net.

The $origin value may be substituted by the @ variable and be used in records where the $origin would otherwise be specified in full (for example, the SOA record).

2 The SOA record denotes the start of authority for the zone sample.net. (as specified by the @). It has no TTL value and cannot be cached. The record contains two domain names; the first is the name of the primary name server for this zone and the second is the mailbox address for the user (JIM) who is responsible for this zone.

The SOA record is split over several lines, the continuation being indicated by the left and right parentheses.

3 This defines the primary name server in this zone.

4 This defines the IP address for host HO5 (the name server).

5 This defines the well-known services running on HO5.

6 This is an information record defining the CPU and operating system for host HO4. Notice that this record has no name in column one. In this case, the name from the previous resource record is used.

7 Two MX records show how host HO1 will receive its mail. The record with the lowest preference value identifies the primary mailbox, in this case HO1 itself. If HO1 is not available for any reason, mail will be delivered to an alternate host, HO2.

8 This is a text record relating to host RO1. These are often used to indicate location information.

9 This record provides the primary (or canonical) name for an alias. If we queried the name server for host2, it would find the CNAME record for host2 pointing at HO2. It would then look up the A record for HO2 and return the address 172.16.1.2.

11.4.9.3 IP Address to Domain Name Mapping

The one common resource record that we did not see in our example was the PTR or *pointer* record. This record is used for mapping addresses to names (the opposite of the A record). While we noted that DNS was established to allow us to use more understandable names rather than addresses, it is also true that a lot of software today actually reverses the process and requires that addresses be mapped onto names. Inetd and rlogin are examples of this. Network management software uses DNS to provide names instead of addresses so that it may provide more easily readable reports.

The Domain Name System provides another part of the name space to offer this service, known as the *in-addr.arpa* zone. Another master file is constructed with the same syntax as the standard DNS master file, but the resource names are IP addresses instead of names. The addresses are also written in reverse order and in-addr.arpa is appended to each. There is one PTR record for each interface on this network and each record can only point to one (canonical) name.

The following is the in-addr.arpa master file for our sample network.

```
;
$origin 16.172.in-addr.arpa.  1
;
@             IN SOA HO5.sample.net. JIM.HO5.sample.net.  (
                     19950517          ;serial number for data
                     10800             ;secondary refreshes every 3 hour
                     3600              ;secondary retries every 1 hour
                     604800            ;data expire after 1 week
                     86400)            ;minimum TTL for data is 1 day
;
@       99999 IN NS  HO5.sample.net.
;
3.2  2  99999 IN PTR HO5.sample.net
;
2.2     99999 IN PTR HO4.sample.net.
1.2     99999 IN PTR HO3.sample.net.
1.1     99999 IN PTR HO1.sample.net.
2.1     99999 IN PTR HO2.sample.net.
1.3     99999 IN PTR HO6.sample.net.
2.3     99999 IN PTR HO7.sample.net.
;
;RO1 and RO2 are routers and each have 2 different IP addresses
;
3.1     99999 IN PTR RO1.sample.net.
4.2     99999 IN PTR RO1.sample.net.
5.2     99999 IN PTR RO2.sample.net.
3.3     99999 IN PTR RO2.sample.net.
;
```

Figure 205. in-addr.arpa Master File on the Name Server HO4

Notes:

1 The $origin statement identifies the origin of the 172.16 network. The address in this statement has the special value in-addr.arpa appended.

2 We only need to show the last part of the address here (in reverse order) as the $origin value will be appended to all domain numbers that do not end in a period. 3.2 will become 3.2.16.172.in-addr.arpa.

Figure 206 shows a further example. The domain is divided into multiple subnets, requiring multiple SOA and $origin records.

```
$origin 10.in-addr.arpa.  3
;
@                          IN  SOA  hutch.secure.itsc.ral.ibm.com.  hutch.(
        95060201 ; Serial number for this data (yymmdd##)
        86400    ; Refresh value for secondary name servers
        300      ; Retry value for secondary name servers
        864000   ; Expire value for secondary name servers
        3600 )   ; Minimum TTL value

@               99999 IN  NS   hutch.secure.itso.ral.ibm.com.
hutch.secure.itso.ral.ibm.com.    IN  A     192.168.1.18
;
14.2.0    4    99999 IN  PTR    rs600014.secure.itso.ral.ibm.com.
25.2.0         99999 IN  PTR    mvs25.secure.itso.ral.ibm.com.
;
$origin 1.168.192.in-addr.arpa.  5
;
@                          IN  SOA  hutch.secure.itsc.ral.ibm.com.  hutch.(
        95060201 ; Serial number for this data (yymmdd##)
        86400    ; Refresh value for secondary name servers
        300      ; Retry value for secondary name servers
        864000   ; Expire value for secondary name servers
        3600 )   ; Minimum TTL value

@               99999 IN  NS   hutch.secure.itso.ral.ibm.com.
hutch.secure.itso.ral.ibm.com.    IN  A     192.168.1.18
;
18        6     99999 IN  PTR  hutch.secure.itso.ral.ibm.com.
```

Figure 206. in-addr.arpa Multiple Subnet Example

Notes:

3 The first $origin statement identifies the origin of the Class A 10 network and is followed by the first SOA record.

4 The addresses here identify the two hosts in the 10.0.2 subnet.

5 The second $origin value identifies the origin of the Class C 192.168.1 network and is followed by the second SOA record.

6 This identifies the address of the host within the 192.168.1 subnet.

11.4.9.4 DNS Zones

We have used the word *zone* on a number of occasions in the last section without explaining its meaning. Divisions in the domain name space can be made between any two adjacent nodes. The group of connected names between those divisions is called a zone. A zone is said to be authoritative for all the names in the connected region. Every zone has at least one node and, consequently, at least one domain name and all the nodes in a zone are connected. This sounds very much like a domain.

However, there is a subtle difference between a zone and a domain. A zone may contain exactly the same domain names and data as a domain, is often the case. If a name server has authority for the whole domain, then the zone will in fact be the same as the domain. As networks grow, it is common that, for the ease of administration, a domain may be divided into subdomains

with the responsibility for these subdomains being delegated to separate parts of an organization or, indeed, to a different organization completely. When this happens, the authority for those subdomains is usually assigned to different name servers. At this point, the zone is no longer the same as the domain. The domain contains all the names and data for all of the subdomains, but the zone will contain only the names and data for which it has been delegated authority.

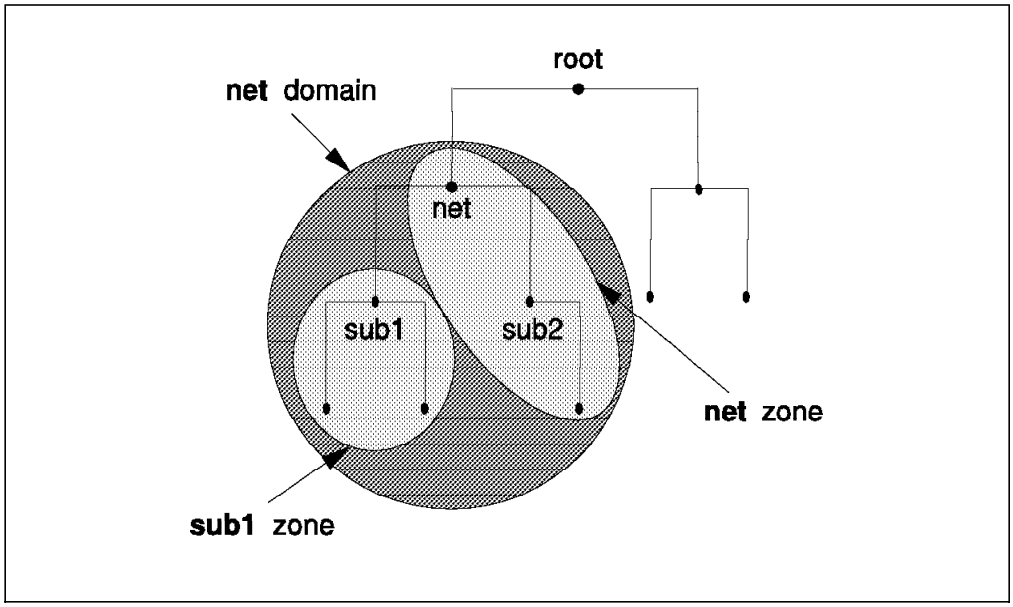


Figure 207. Domains and Zones

Figure 207 illustrates the difference between a zone and a domain. The net domain contains names and data for the net domain, the sub1 domain and the sub2 domain. (Sub1 and sub2 are both subdomains of the net domain). However, only domain sub1 has been delegated the authority for its resources and hence has its own zone, the sub1 zone. The sub2 domain is still under the authority of the net zone.

11.4.9.5 Name Servers

The second component of the Domain Name System is the name server. Name servers are the repositories for all of the information that makes up the domain name space. Originally, there was a single name server, operated by the NIC, which held the single HOSTS.TXT file. The concept of the hierarchical name space has meant that a single name server would be impractical. There are now nine *root name servers* with responsibility for the top-level domains. The name space is then divided into zones, as we have already discussed, and these zones are distributed among the name servers such that each name server will have authority over just a small section of the name space. This division is frequently based on organizational boundaries, with freedom to subdivide at will. A name server may, and often will, support more than one zone, and a single zone may be served by more than one name server.

Name servers come in the following three types:

- Primary name server - This maintains the zone data for the zones it has authority over. Queries for this data will be answered with information from files kept on this name server.

- Secondary name server - This has authority over a zone but does not maintain the data on its own disks. The zone data is copied from the primary name server database when the servers are started. This is known as a *zone transfer*. The secondary then contacts the primary at regular intervals for updates.
- Caching-only name server - This server has no authority over any zones and contains only records pointing to other (primary or secondary) name servers. Data is kept in a cache for future use and discarded after a time-to-live value expires.

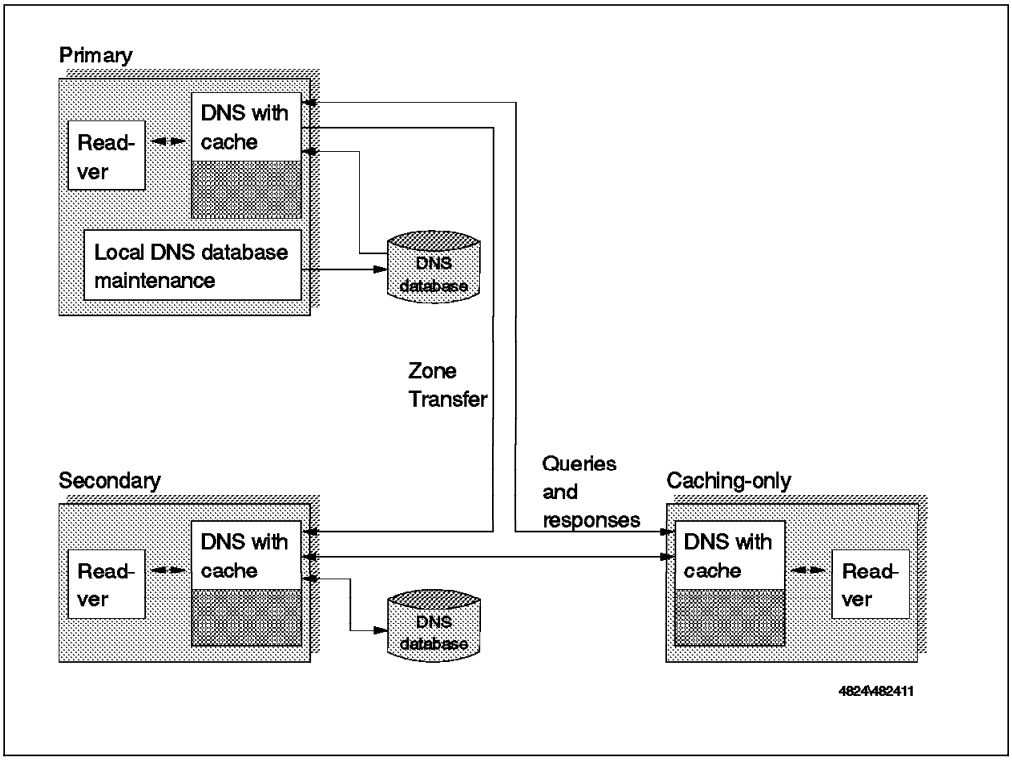


Figure 208. Name Server Categories

The main function of the name server is to answer standard queries from clients. These queries flow in DNS messages and identify the type of information that the client wants from the database and the host in question. The name server can answer queries in a number of ways depending on the mode of operation of the client and server.

- *Recursive mode* - When a client makes a recursive query for information about a specified domain name, the name server will respond either with the required information or with an error, such as the domain name does not exist (name error) or there is no information of the requested type. If the name server does not have authority over the domain name in the query, it will send its own queries to other name servers to find the answer. These name servers are pointed to by the additional NS resource records in the database.

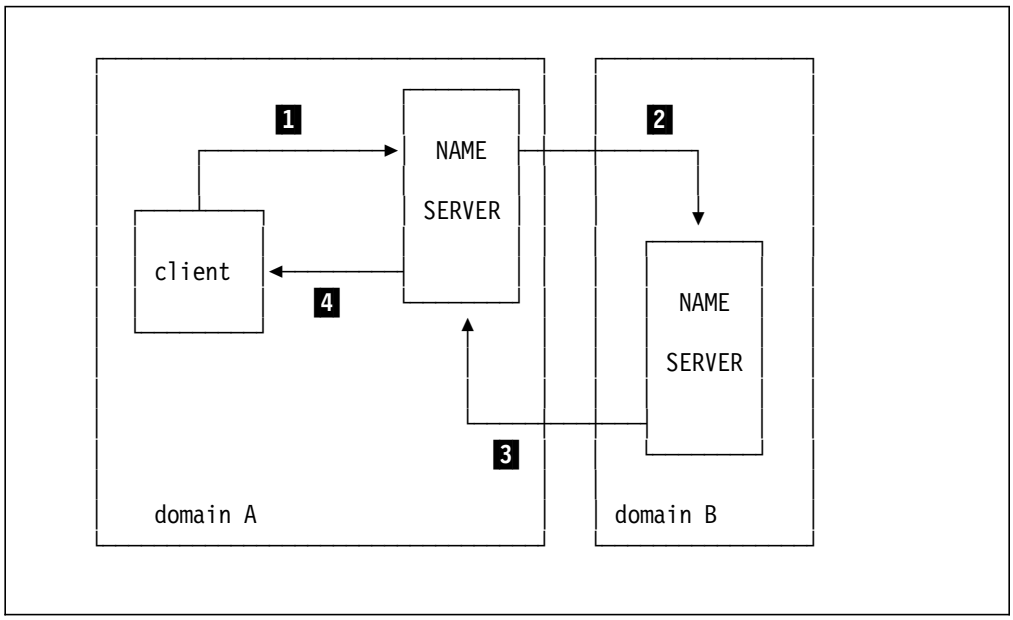


Figure 209. Recursive Mode

Notes:

1 The client in domain A sends a simple query to its name server asking for the address of a host in domain B.

2 The specified name server does not have authority over domain B and has no record of the host. The name server has an NS resource record pointing to an authoritative name server for domain B and so it sends a query to that name server asking for the address of the host.

3 The name server in domain B returns the address of the host to the name server in domain A.

4 The name server in domain A returns the address of the host to the client.

- *Nonrecursive or Iterative mode* - In this case, when a client makes a query, the name server has an extra option. It will return the information if it has it. If not, rather than ask other name servers if they have the data, it will respond to the query with the names and addresses of other name servers for the client to try next.

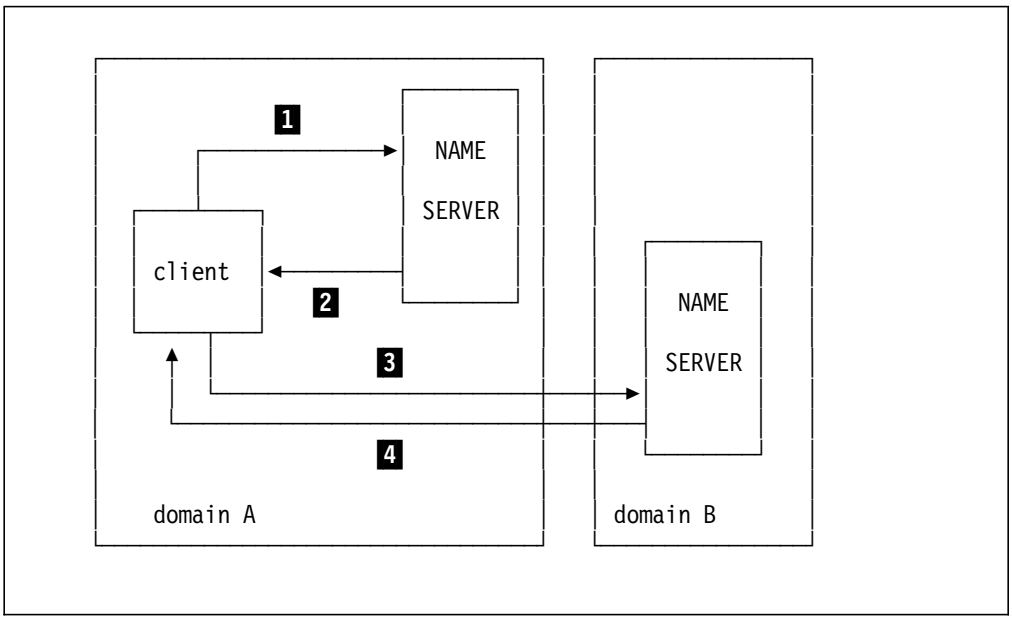


Figure 210. Nonrecursive or Iterative Mode

Notes:

1 The client in domain A sends a simple query to its name server asking for the address of a host in domain B.

2 The specified name server does not have authority over domain B and has no record of the host. The name server has an NS resource record pointing to an authoritative name server for domain B. But rather than send its own query to that name server, it responds negatively to the clients query and gives the client the address of the name server in domain B.

3 The client sends a second query, this time to the name server in domain B.

4 The name server in domain B returns the address of the host to the client.

11.4.9.6 Resolvers

The resolvers are the third component of the Domain Name System. These are the clients making queries to the name servers on behalf of programs running on the host. These user programs make system or subroutine calls to the resolver, requesting information from the name server. The resolver, which runs on the same host as the user program, will transform the request into a search specification for resource records located (hopefully) somewhere in the domain name space. The request is then sent as a query to a name server, which will respond with the desired information to the resolver. This information is then returned to the user program in a format compatible with the local host′s data formats.

What exactly does the resolver have to do for the client program? There are typically three functions that need to be performed:

1. Host name to host address translation

The client program (for example, FTP or Telnet) will provide a character string representing a host name. This will either be a fully qualified domain name (host.net.com.) or a simple unqualified host name. Let us use HO4 from our previous example. If the name is unqualified, the resolver code will append a domain origin name (in our case sample.net.) to the name before passing it to the server. This domain origin name is one of four parameters that are configured on every IP host:

IP address of the host

Host name

Domain origin name - The domain to which this host belongs

IP address of the name server(s) being used

The resolver then translates this request into a query for address (type A) resource records and passes it to the specified name server. The server will return one or more 32-bit IP addresses.

2. Host address to host name translation

 Presented with a 32-bit IP address from the client program (perhaps SNMP), the resolver will query the name server for a character string representing the name of the host in question. This time the query is for PTR-type resource records from the in-addr.arpa name space. The resolver will reverse the IP address and append the special characters in-addr.arpa before passing the query to the name server.

3. General lookup function

 This function allows the resolver to make general queries to the name server requesting *all* matching resource records based on the name, class and type specified in the query.

There are two types of resolvers, both making use of the routines gethostbyname() for name to address translation and gethostbyaddr() for address to name translation. The first, known as a *full resolver*, is a program distinct from the client user program. The full resolver has a set of default name servers it knows about. It may also have a cache to retain responses from the name server for later use.

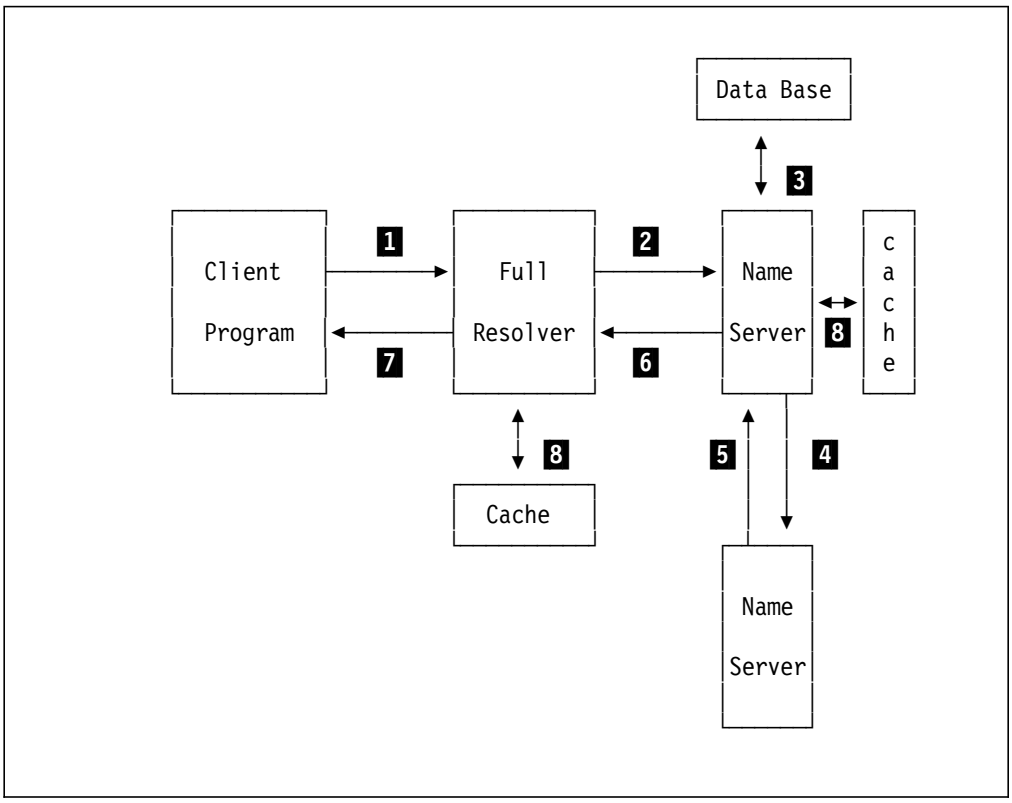


Figure 211. A DNS Full Resolver

Notes:

1 The user program makes a call to the resolver.

2 The resolver translates the call into a resource record query and passes it to its default name server.

3 The name server will attempt to resolve the query from its own database. Assume that this is the first query and there is nothing in the cache.

4 If unable to locate the requested records in its own database, the name server will pass its own query to other name servers that it knows (if recursive mode is being used).

5 The remote name servers eventually reply with the required information.

6 The local name server passes the information back to the resolver.

7 The resolver translates the resource records into local file format and returns the call to the user program.

8 Both the resolver and the name server will update their caches with the information.

The second, and possibly more common, type of resolver is the *stub resolver*. This is merely a routine or routines which are linked to the user

program. The stub resolver will perform the same function as the full resolver but generally does not keep a cache.

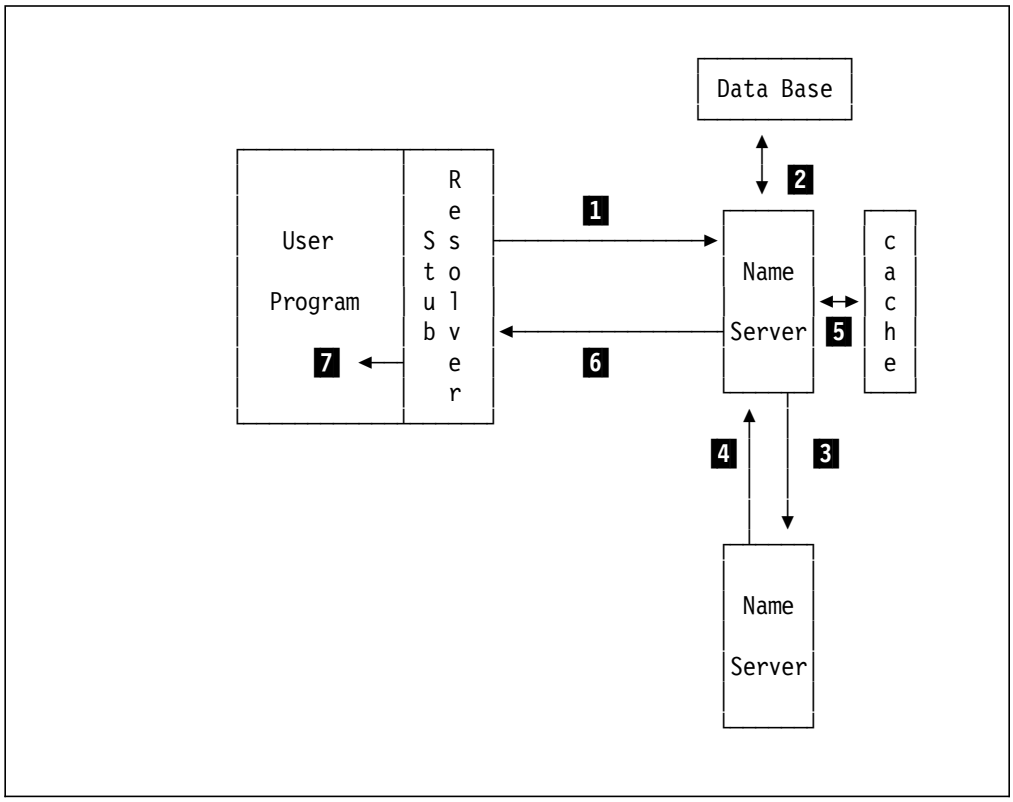


Figure 212. A DNS Stub Resolver

Notes:

1 The user program invokes the stub resolver routines; the resolver creates an RR query and passes it to its default name server.

2 The name server will attempt to resolve the query from its own database. Assume that this is the first query and there is nothing in the cache.

3 If unable to locate the requested records in its own database, the name server will pass its own query to other name servers that it knows (if recursive mode is being used).

4 The remote name servers eventually reply with the required information.

5 The name server will update its cache with the information.

6 The local name server passes the information back to the resolver.

7 The resolver translates the resource records into local file format and returns the call to the user program.

11.4.10 Mail Support

We stated earlier that the Domain Name System not only includes functions for name to address translation and vice versa but also provides a repository for useful information about the nodes in the name space. One such example of this added value is the support that DNS provides for mail services.

DNS has defined a standard for mapping mailbox names into domain names using MX (mail exchange) resource records. It also defines the way in which these records are used to provide mail routing within the Internet. The standards define a mailbox name in the form <local-part>@<mail-domain>. For the exact syntax of this form, please refer to RFC 822 - Standard for the Format of ARPA Internet Text Messages. DNS encodes the <local-part> as a single label. Any special characters in the original character string can be preserved in the DNS master file label by using backslash quoting. For example, the name Mail.server would be coded as Mail\.server. The <mail-domain> is simply encoded as a domain name and appended to the mailbox label. Thus, the mailbox name Mail.server@sample.net. would have a DNS MX record name of Mail\.server.sample.net.

The DNS MX record actually has two values in the rdata section. The one we have just seen previously is the name of the mailbox host. The other is an unsigned 16-bit integer which acts as a preference value. This is used to indicate a priority to the MX records if there is more than one for this domain name. The lower the preference value, the higher the priority. The following example illustrates this:

```
sample.net    MX  5   Mail\.server.sample.net.
              MX  10  Mailbox.sample.net.
```

We have two mailboxes defined for the sample.net. domain. The first mailbox Mail\.server has a preference value of 5 and so is higher in priority to the second mailbox (Mailbox), which has a preference value of 10. If the mail system has mail for user@sample.net., then it will use the MX records for the sample.net. mail domain, as seen previously, and will attempt to deliver the mail to the mailbox with the lowest preference value (in this case Mail\.server.sample.net.). If this mailbox is unavailable, the mail system will try Mailbox.sample.net.

11.4.11 DNS Design Requirements

We have spoken at some length on the technicalities of the Domain Name System. We shall restrict ourselves in the rest of this chapter to the design considerations necessary for the implementation of a Domain Name System within your network.

11.4.11.1 Do I Need DNS?

The first question you are bound to ask is whether you really need a Domain Name System. We began this book by advising you that we are aiming not at the casual user of the Internet, perhaps dialing in from a home PC, but instead at the organization that needs to build an internetwork of its own. The answer must therefore, in part, be based on the size of the network in question. DNS began life as a single flat name space and that scheme still lives on today in the form of the ETC\HOSTS file (sometimes known as the HOSTS.LOCAL file). However, instead of being one enormous file on a single server, the ETC\HOSTS file is a small local file on each host, identifying

frequently used local host names and mapping them to their IP addresses. In a small network environment, perhaps a single LAN, the ETC\HOSTS file is often sufficient for your name resolution requirements. Put together several LANs and add a router or two, and the resulting network is really too large for the ETC\HOSTS file to handle on its own. At this point you will need to implement DNS.

The size of the network is not the only dependency for DNS. Take another look at the added value you get with DNS and you will see that functions such as e-mail depend heavily on the name server process. Most of the other Internet services and standard functions, such as remote terminal access and file transfer, are made more efficient to users by using names rather than addresses, so name resolution is a major requirement.

Domain Administration

Let us assume that you have decided to implement DNS. The next question you ask is who is going to set up and run the domain. Again, the answer may depend on the size of the network. A reasonably small network may (and probably will) be able to take advantage of the services offered by its IP service provider, perhaps becoming part of the service provider′s domain. As the network grows, you will undoubtedly be seeking your own identity and wish to establish your own domain. But again, you may not need to do all the work yourself. The IP service provider may be happy to set up your domain and administer it at a price.

However, if you decide to administer a domain yourself, the key requirement is that you have a designated manager for supervising that domain′s name space. This person (or persons) will be the technical and administrative contacts for the domain. They are the trustees of the domain and have a duty to serve the network community.

Registering a Domain

The third question is about how to set up your own domain. There are several parts to this issue and it is not as easy as it may seem at first. We start by making a few assumptions:

1. The network is large enough to require its own domain.
2. The users of the network require access to other networks (such as the Internet).
3. Functions such as e-mail will be utilized within the network.

Let′s begin with the name. Setting up a domain implies that you will have a name for that domain. After all, that′s what it′s all about. This name will have to be registered with one of the Internet authorities. The IANA has the overall responsibility for the domain name space on the Internet and for delegation of the top-level domains (TLDs). The day-to-day administration of the Domain Name System is performed by the Internet Registry (IR), which is currently the InterNIC. As with the IP network numbers, growth in Internet activity has led to a further delegation of authority for the domain name space. Two regional bodies, RIPE NCC (Reseaux IP Europeans Network Coordination Centre) and APNIC (Asia Pacific Network Information Centre) now handle the domain name space requirements for Europe and Asia Pacific, respectively. Requests for registration of domain names should be

sent to the appropriate authority (InterNIC in the US or to a delegated regional/national authority).

After submitting a registration request, InterNIC or the regional or national authority will search for any other occurrence of the name selected and then register it. This is a paid process and the turnaround time is not under your control. There can be a delay of some weeks, registering domain names under top-level domains, such as .com, controlled by InterNIC in the US.

Before making your request for a domain name, you will need to know where you fit within the domain name space. The place to start is the top-level domain. Which TLD do you fit in? If you are outside the United States, you will need to find out if your country has a TLD registered. The second-level domain structure will vary from country to country, but often takes the form of co or com for commercial companies, ac for academic bodies, go for government organizations and re for research groups. If in doubt, talk to your IP service provider or directly to the regional Internet authorities. Depending on the service being offered by your service provider, you may become a subdomain of their own domain (for example, your-domain.vendor.co.uk.).

There are rules governing the name you choose for your domain. The general rule of thumb is to keep the name short and simple (most domain names are actually under 12 characters) but names of three or less characters are usually reserved. However, remember that using a name for your domain and registering that name with the relavent Internet authority does not give you any legal rights to that name. The IANA will not usually get involved in local disputes over a name, but there is a committee, the Internet DNS Names Review Board (INRB), which may act as a review panel in some cases.

When you establish a domain for your organization you are being delegated the responsibility for a new branch of the domain name space tree structure. This delegation is being done by your parent domain, and it follows that you will have to register your domain with them as well. Some countries put organizations directly below the country TLD (Canada and France, for example) but others place your domain as a third-level domain or lower. For example, company.co.uk. would place the domain for company below the second-level domain co for commercial organizations in the UK. You can apply for your domain to be under one of the generic TLDs (for example, in the case that multinational companies may use .com as their parent domain).

Establishing Name Servers

When registering for a domain name, it is wise to read all the small print. The InterNIC registration templates have some words to say about your name servers. To begin with, note that they stipulate that a domain must provide at least two independent name servers, one primary and one secondary, which should be in different physical locations and on different networks (if possible). They also require that these name servers be active and responsive to DNS queries before you submit your application. The first part of that directive is good advice and should be adhered to (if possible). Creating a domain of your own and making name resolution services available to your network community will quickly turn the name server(s) into a critical service requiring 24 hours a day, seven days a week availability.

As a single point of failure, it makes sense to have a secondary server should the inevitable happen. In practice, as the network grows it is probable that you will require more than two name servers to cope with the growth in queries. Also, the growth in the network will probably result in a delegation of responsibility for parts of the network; this is an ideal time to create subdomains and to establish name servers with authority for the new zones that encompass them.

The location of your name servers is also important, not only in terms of physical location but also in the choice of the hosts that will provide the service. A name server must be well connected. It will need to be accessible by all the hosts on your network, so we recommend you don't place it in some remote corner of the network on the end of your slowest line. Place the primary name server at the hub of your network within easy access from all parts of the network and from your Internet access point.

Try to utilize multihomed hosts that can directly serve more than one part of your network. Also, unless you are going to dedicate a host to the name server function, choose a multiuser host rather than single user systems. If a large portion of your users reside on an MVS mainframe, why make them traverse your network every time they want to query a name? By running the name server on the mainframe you will reduce the time delays for most of your users with a corresponding reduction in network traffic. It is also beneficial to run multiple secondary name servers, perhaps each serving one or more subnets in your network. The major administrative burden is carried on the primary name server with each secondary obtaining its information by zone transfers. The correct placing of these secondaries can also reduce total network load, with less frequent zone transfers from the primary taking the place of large numbers of frequent queries to the primary.

11.4.11.2 DNS Security

The problem is that with DNS we are aiming to provide a name service to actually allow people in our network to be found. We must therefore adopt a special technique when installing a name server in relation to a firewall. This obviously has implications for e-mail as well.

The goal of this scheme is to provide a full domain name service to hosts inside the secure network while only providing information about the firewall itself to the outside world. Let us assume you have already set up one or more name servers within your network. These will remain virtually unchanged and will serve your secure hosts, giving them information about your secure network. You will need to set up a new name server on the firewall. This is often provided as a feature of the firewall implementation. The firewall name server will respond to queries from the outside only with information about the firewall address itself. When a host in your secure network makes a query about a host in the nonsecure network, the name server will forward the query to the firewall name server. The firewall name server will, in turn, refer the query to a name server in the nonsecure network, probably the one provided by your Internet service provider.

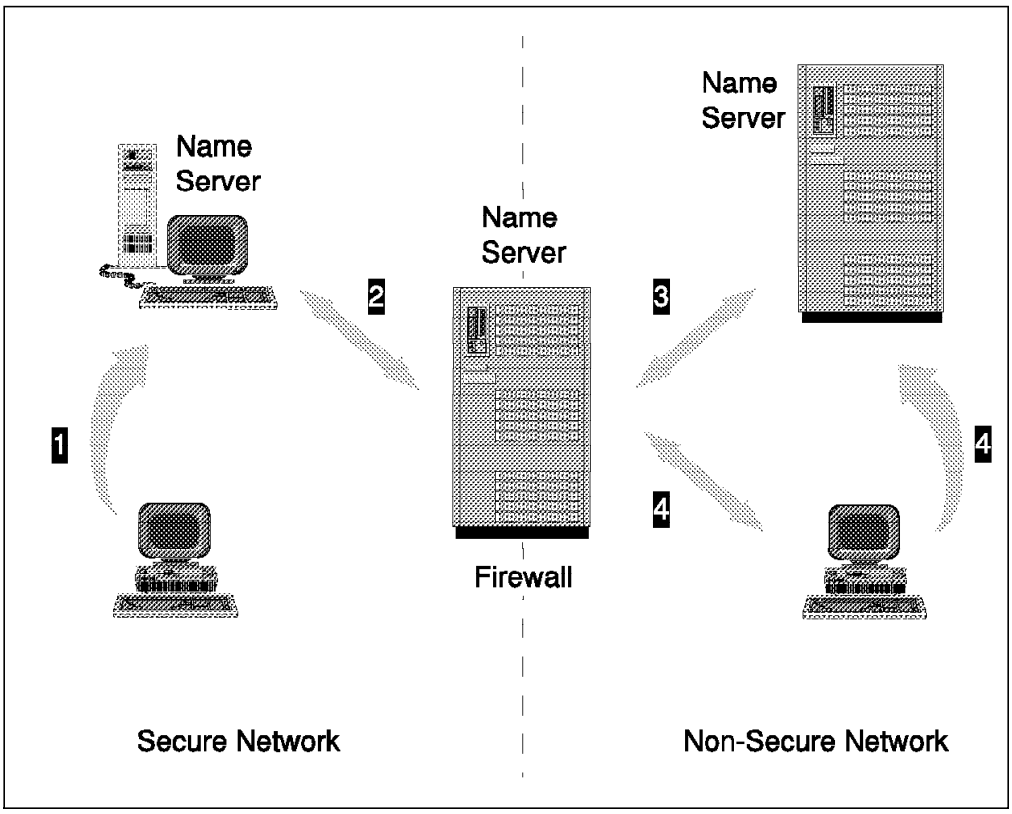


Figure 213. DNS Coexistence with Firewalls

Notes:

1 Hosts inside the secure network make their normal requests to an internal name server. Local domain names are returned directly.

2 Queries for names in external domains are passed by the internal name server to the firewall name server.

3 The firewall name server will pass the queries to an external name server, and the responses will follow the same route back to the original internal host.

4 Queries from external hosts will be directed either through an external name server or directly at the firewall name server, but in either case the firewall name server will respond with a *restricted* answer.

A similar process applies to electronic mail passing through the firewall. One way to overcome the problem is to employ a mail forwarding service on the firewall. This will act as a relay for the secure mail server inside the secure network. External hosts will direct their mail at user@firewall.company.com or user@company.com depending on where the domain begins. Both the secure mail server and the mail forwarder on the firewall must be configured as Relay Hosts (DR entry) to allow mail headers to be rewritten and mail not destined for the local host to be routed through the firewall.

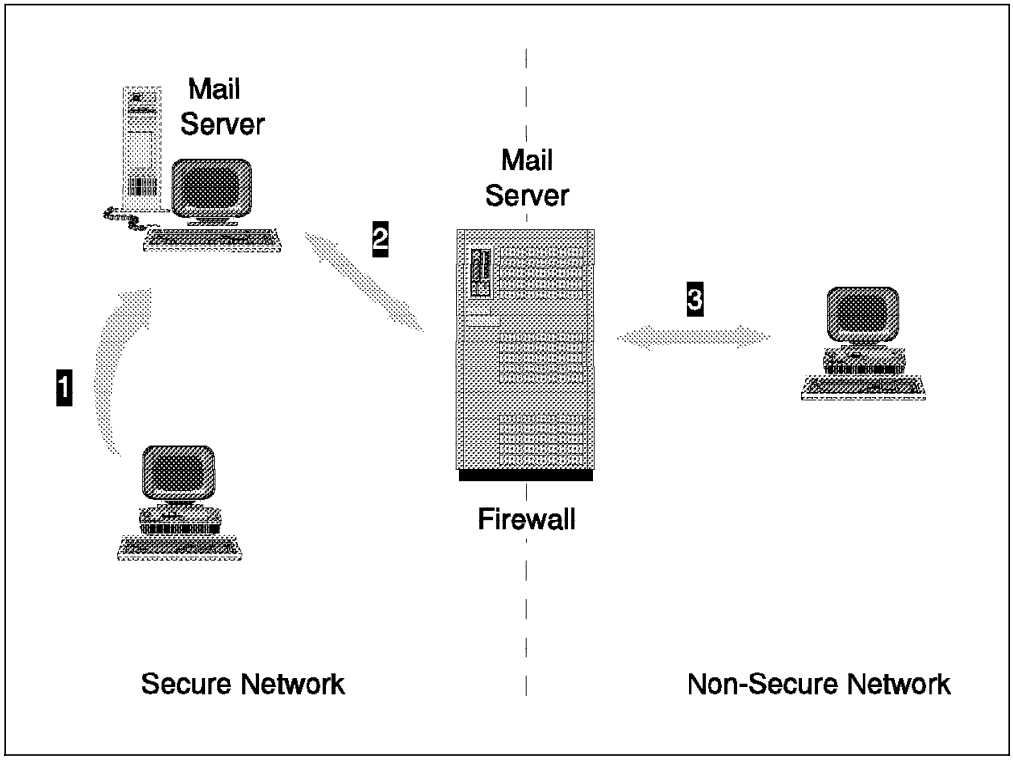


Figure 214. DNS and E-Mail with Firewalls

Notes:

1 Internal hosts use the secure mail server to deliver mail within the secure network (or deliver directly themselves).

2 Mail destined for external users is passed to the secure mail server for outbound relay to the firewall mail server.

3 The firewall routes mail to the outside world. Inbound mail cannot be directly delivered to internal users but must be relayed through the firewall to the secure mail server, which has ultimate responsibility for delivery of the mail.

For further information about IP network design, refer to *The Basics of IP Network Design*, SG24-2580-00.

11.4.12 Web Server Softwares

When building a content service (for example, content hosting), you need to have a software installed that supports all the customers′ environments. This software must run on a server machine that must be able to host the customers′ Web pages.

The Web server softwares are responsible for conducting secure electronic commerce and communications on the Internet and other TCP/IP-based networks. Capacity, availability and reliability are its main tasks in order to support this special environment. The speed, the design and the number of features are some of the items that must be considered before choosing which software you must buy.

After learning how to submit a request to InterNIC or to another regional/national authority in order to get a new range of IP addresses and a domain name to your Web site or to your customers, you need to know further details about the installation, configuration and use of a Web server SW.

In the remainder of this chapter, you are going to see some considerations about the use of multiple IP addresses, the setup of your Web server SW, how to work with it, how to set up an FTP anonymous area and all the management and operational issues you must be aware of.

The example we adopted uses the Netscape Commerce Server in order to show how to configure, run and access a Web server SW for the AIX (IBM UNIX) environment.

For further information about Web server SWs and their hardware and softwares requirements, refer to Chapter 3, “Additional IBM Software Solution” on page 155.

11.4.13 Multiple IP Addresses

You can use just one host to be your server with one IP address associated to it or you can have multiple IP addresses to the same host.

In the second case you are going to need a reduced number of workstations to create your content hosting services. That is, you are going to install more than one customer and their logical servers by each server workstation.

The number of logical servers/customers in each workstation depends on the maximum number of processes your Web server software is supposed to attend at the same time and the use of memory, CPU time and disk space. This item depends on the amount to be allocated to each customer.

In order to create multiple IP addresses to your workstation, you′re supposed to use an alias to your host IP address.

In the UNIX environments you can use the ifconfig command to create IP address aliases to your host. For instance, let′s consider you are using a token-ring network adapter tr0, and the new IP address you want to create as an alias is 9.24.104.237, the command should be:

```
ifconfig tr0 9.24.104.237 alias
```

You can use this command whenever you want to create an IP address alias until the customers′ limit number on your Web server.

If you are going to have just one customer/server in your host, then it′s not necessary to use this command.

11.4.14 Setting up the Netscape Commerce Server for AIX

Once you have already created the alias IP address or decided to use your own host IP address, log in to the server host as ROOT, ADMIN or an equivalent user ID.

Before setting up your Web server software, you must create a customer user ID and password and a respective file system according to your

customers disk space size definition based on your content services marketing offering.

After creating the respective ID and file system, change the directory to the .../https/install directory, wherever it is and begin to install your logical Netscape server using the ns-setup program.

For example, type:

```
cd .../https/install
```

Type **./ns-setup** and press Enter to begin server installation.

Indentify your machine's full name in the next screen.

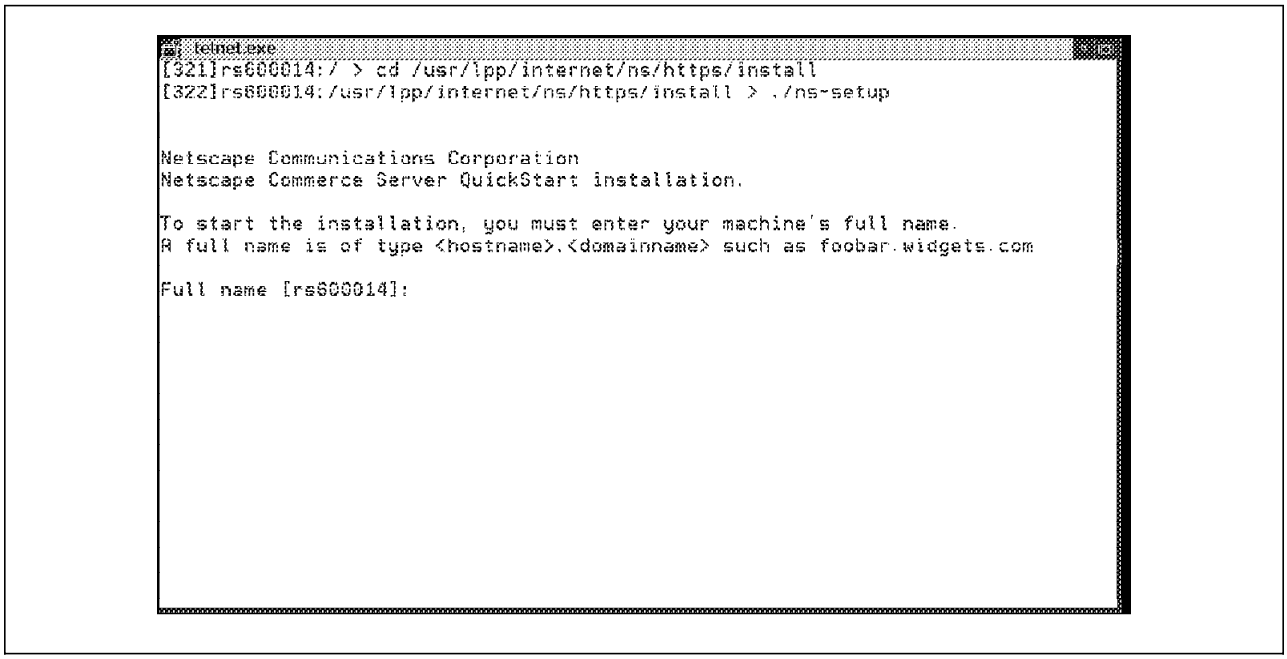


Figure 215. Web Server Installation/Configuration

Enter the name of the browser you'll use to see the forms.

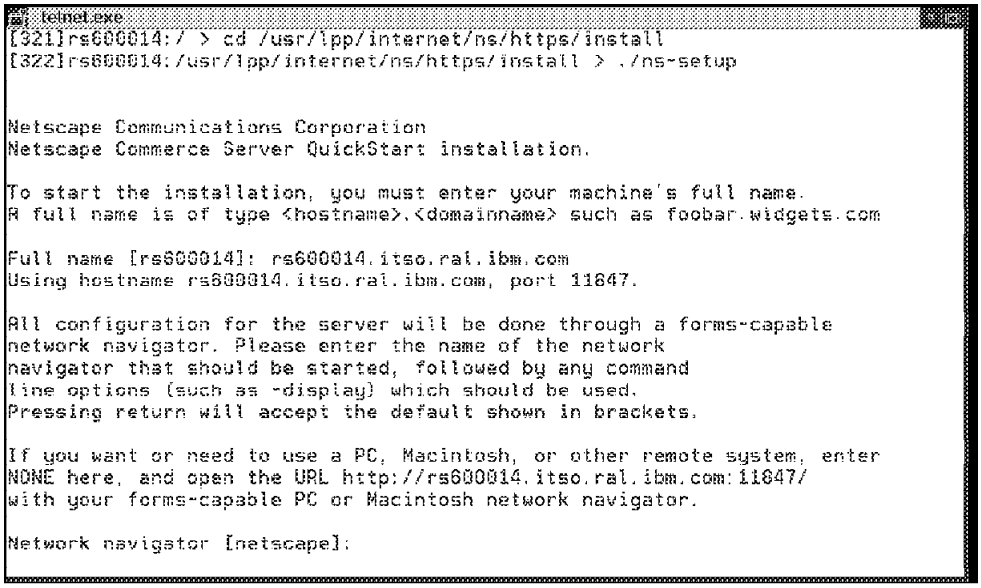


Figure 216. Web Server Installation/Configuration

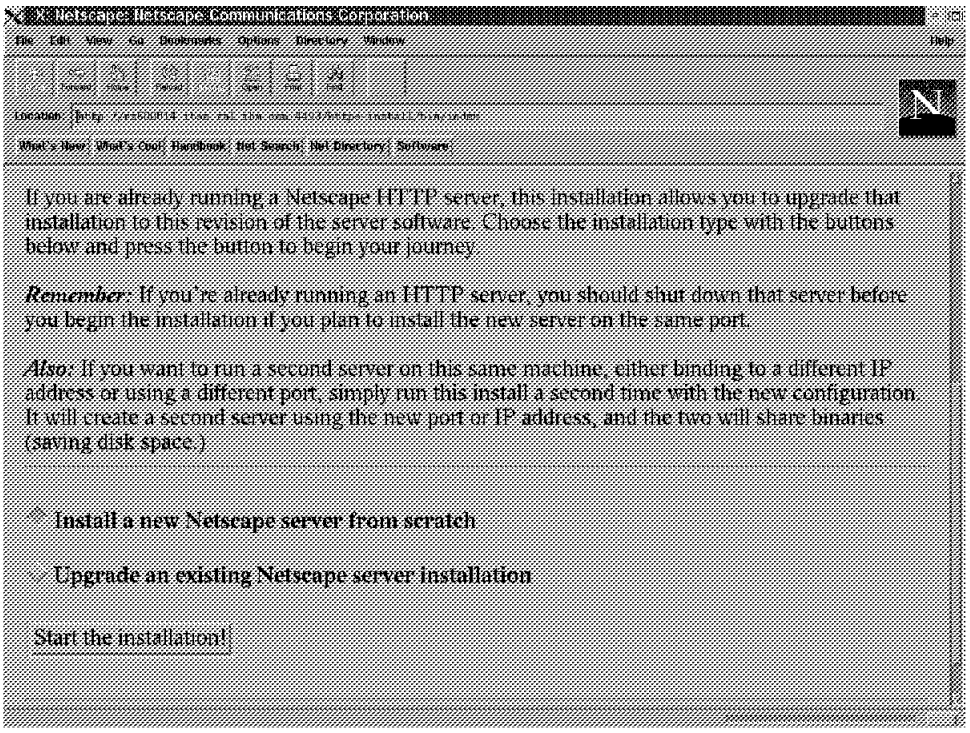


Figure 217. Web Server Installation/Configuration

Select the option **Install a new Netscape server from scratch** and click on the ***Start the installation!*** button.

If you are already running a Netscape HTTP server, this installation allows you to upgrade that installation to this revision of the server software.

Remember, if you are already running an HTTP server, you should shut down that server before you begin the installation if you plan to install the new server on the same port.

Also, if you want to run a second server on your machine, either binding to a different address or using a different port, simply run this installation a second time with the new configuration. It will create a second server using the new port or IP address and the two will share binaries (saving disk space).

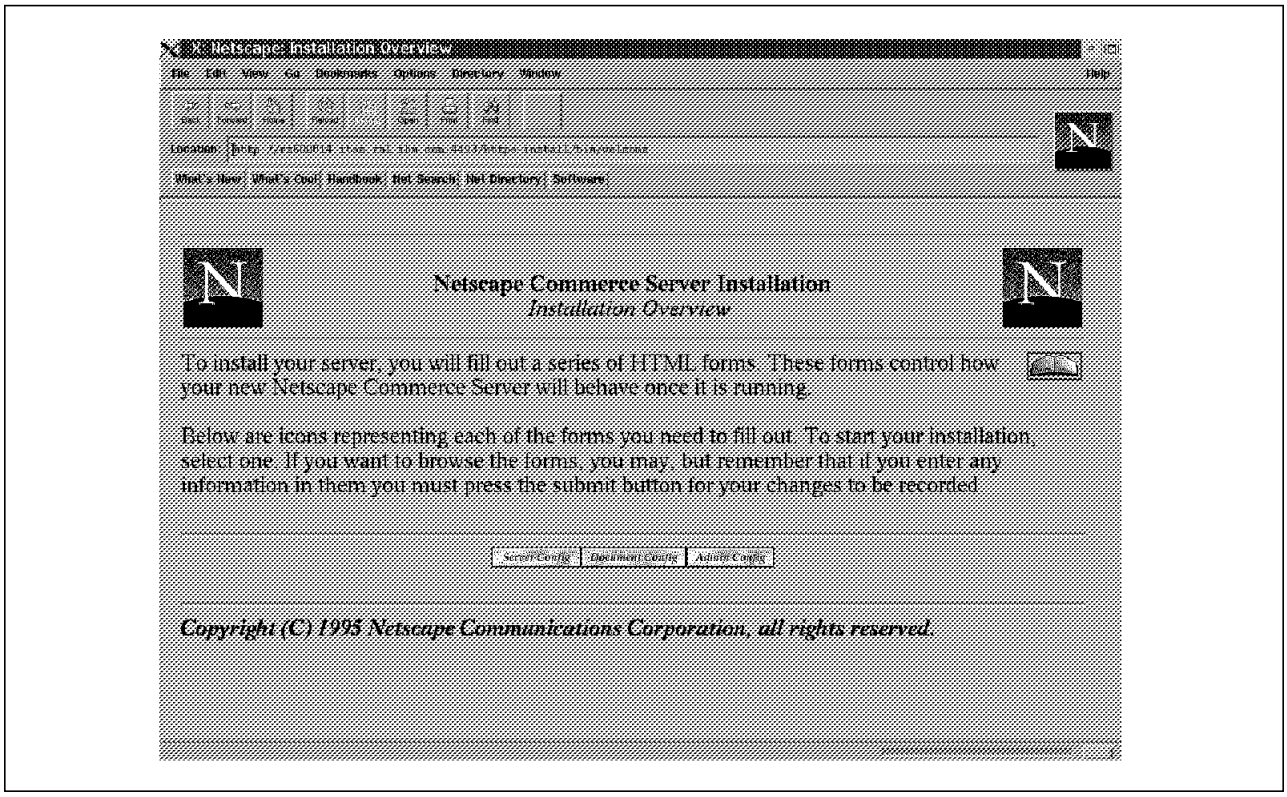


Figure 218. Web Server Installation/Configuration

Select the option you want to configure first: Server Config, Document Config or Admin Config and click on the corresponding button.

Each one of the options shows you a different form to fill out. These forms will implement your server installation/configuration.

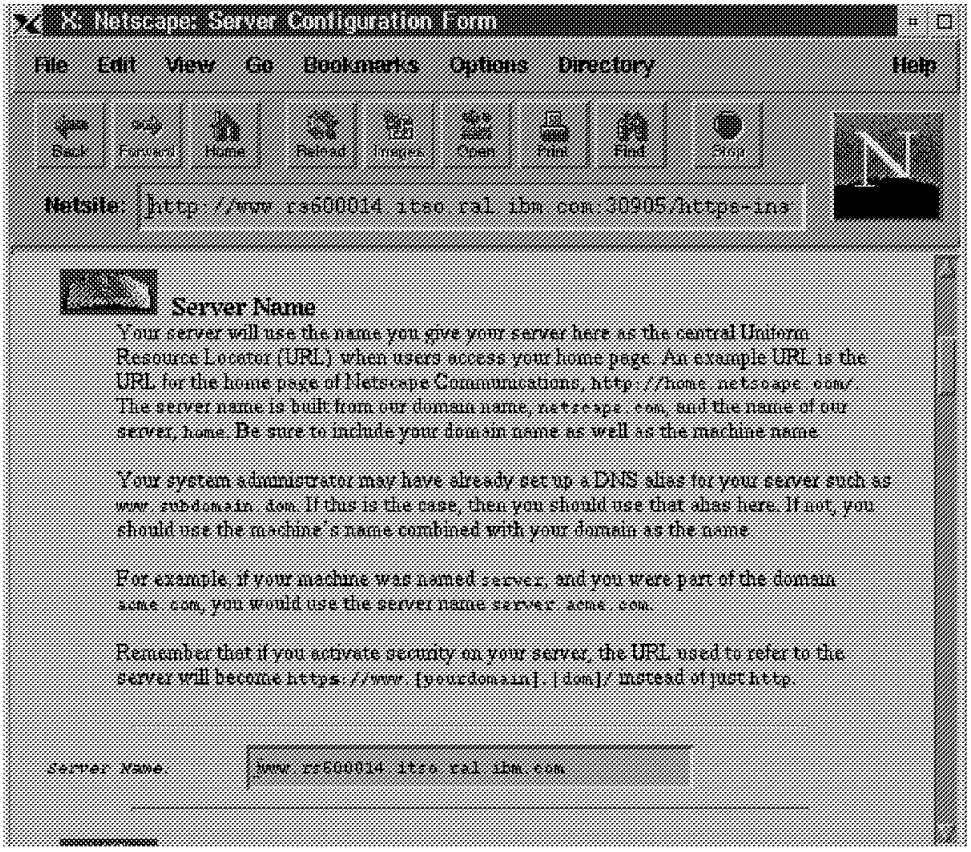


Figure 219. Web Server Installation/Configuration

Your server will use the name you give your server here as the central Uniform Resource Locator (URL) when users access your home page. An example URL is the URL for the home page of IBM, http://www.ibm.com. The server name is built from our domain name, ibm.com, and the name of our server, www. Be sure to include your domain name as well as the machine name.

Your system administrator may have already set up a DNS alias for your server such as www.subdomain.dom. If this is the case, then you should use that alias here. If not, you should use the machine's name combined with your domain as the name.

Remember that if you activate security on your server, the URL used to refer to the server will become http**s**://www.yourdomain.dom/ instead of just http.

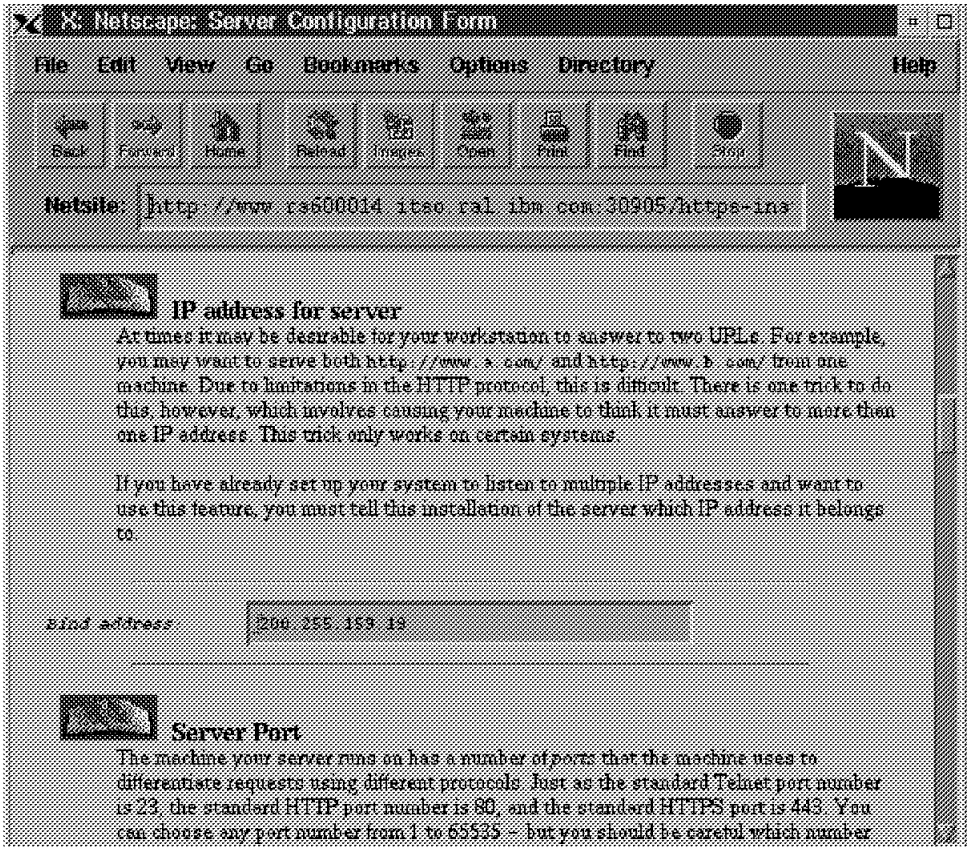


Figure 220. Web Server Installation/Configuration

At times it may be desirable for your workstation to answer to two URLs. For example, you may want to serve both http://www.a.com/ and http://www.b.com/ from one machine. Due to limitations in the HTTP protocol, this is difficult. However, there is one trick to do this, which involves causing your machine to think it must answer to more than one IP address. This trick only works on certain systems.

If you have already set up your system to listen to multiple IP addresses and want to use this feature, you must tell this installation of the server which IP address it belongs to.

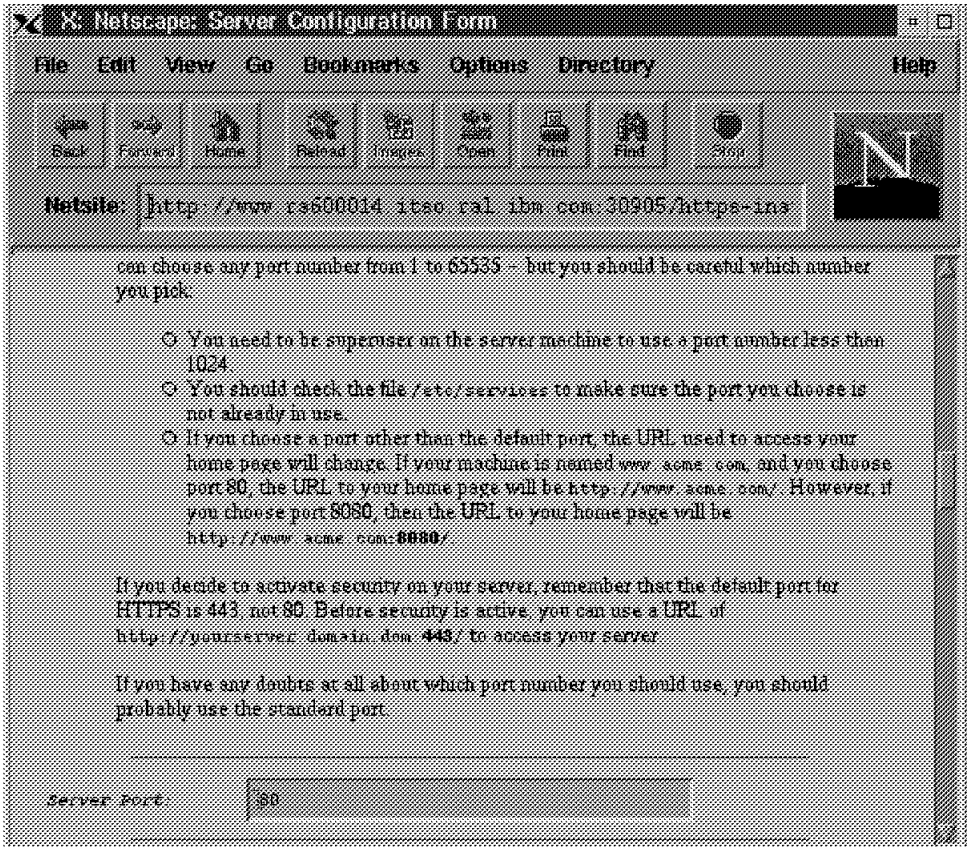


Figure 221. Web Server Installation/Configuration

The machine your server runs on has a number of ports that the machine uses to differentiate requests using different protocols. Just as the standard Telnet port number is 23, the standard HTTP port number is 80, and the standard HTTPS port is 443. You can choose any port number from 1 to 65535, but you should be careful which number you pick:

- You need to be superuser on the server machine to use a port number less than 1024.
- You should check the file /etc/services to make sure the port you choose is not already in use.
- If you choose a port other than the default port, the URL used to access your home page will change. If your machine is named www.acme.com, and you choose port 80, the URL to your home page will be http://www.acme.com/. However, if you choose port 8080, then the URL to your home page will be http://www.acme.com:8080.

If you decide to activate security on your server, remember that the default port for HTTPS is 443, not 80. Before security is active, you can use a URL of http://yourserver.domain.dom:443/ to access your server.

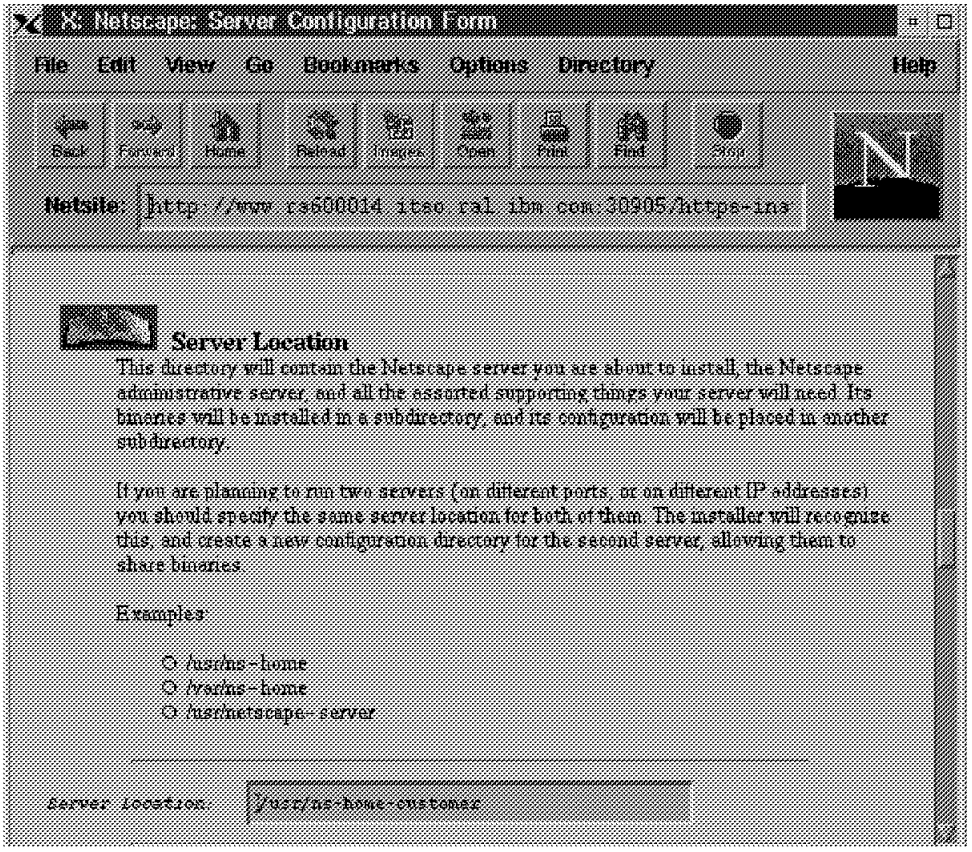


Figure 222. Web Server Installation/Configuration

The server location directory will contain the Netscape server you are about to install, the Netscape administrative server, and all the assorted supporting things your server will need. Its binaries will be installed in a subdirectory, and its configuration will be placed in another subdirectory.

If you are planning to run two servers (on different ports or on different IP addresses), you should specify the same server location for both of them. The installer will recognize this and create a new configuration directory for the second server, allowing them to share binaries.

Examples:

- /usr/ns-home
- /var/ns-home
- /usr/ns-customername

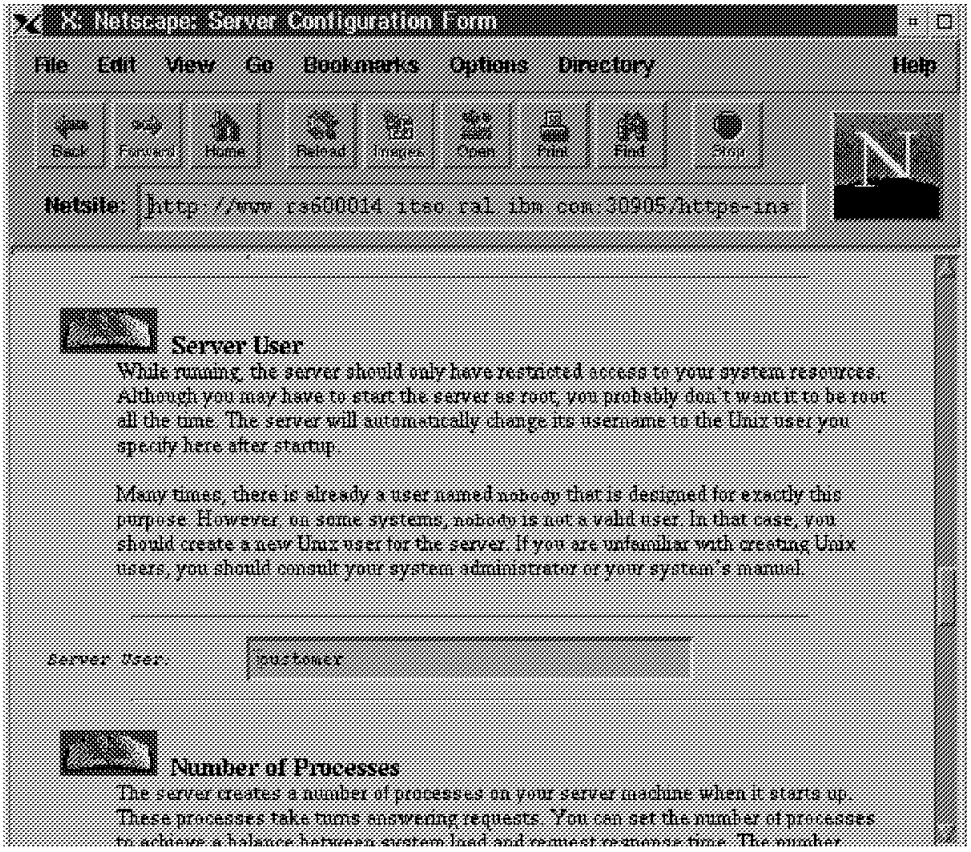


Figure 223. Web Server Installation/Configuration

While running, the server should only have restricted access to your system resources. Although you may have to start the server as root, you probably don't want it to be a root all the time. The server will automatically change its user name to the UNIX user you specify here after startup.

Many times, there is already a user named nobody that is designed for exactly this purpose. However, on some systems, nobody is not a valid user. In that case, you should create a new UNIX user for the server. If you are unfamiliar with creating UNIX users, you should consult your system administrator or your system's manual.

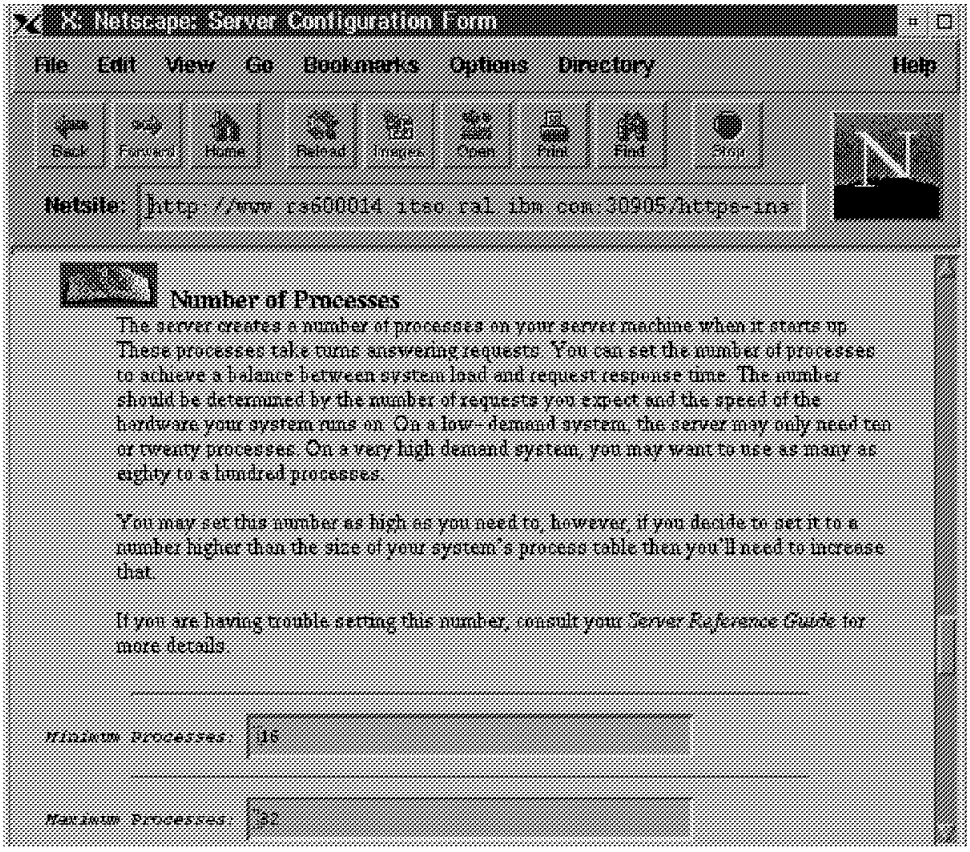


Figure 224. Web Server Installation/Configuration

The server creates a number of processes on your server machine when it starts up. These processes take turns answering requests. You can set the number of processes to achieve a balance between the system load and request response time. The number should be determined by the number of requests you expect and the speed of the hardware your system runs on. On a low-demand system, the server may only need ten or twenty processes. On a very high-demand system, you may want to use as many as eighty to one hundred processes.

You may set this number as high as you need to. However, if you decide to set it to a number higher than the size of your system's process table, then you'll need to increase the size of your table.

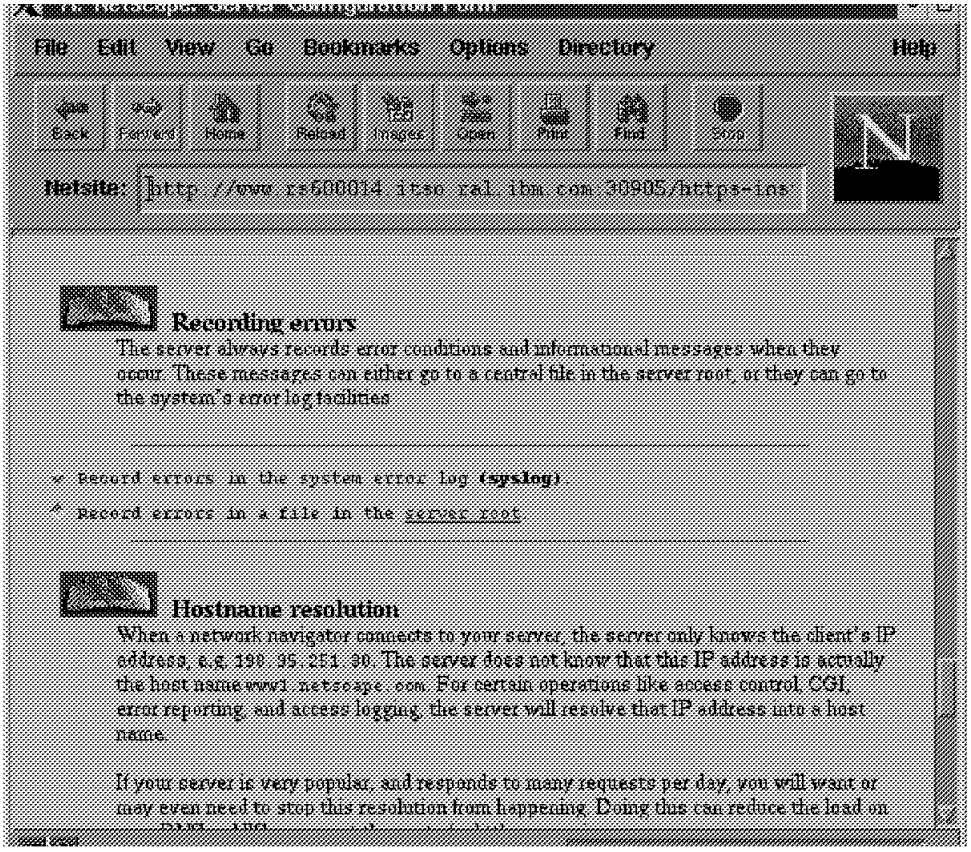


Figure 225. Web Server Installation/Configuration

The server always records error conditions and informational messages when they occur. These messages can either go to a central file in the server root, or they can go to the system's error log facilities.

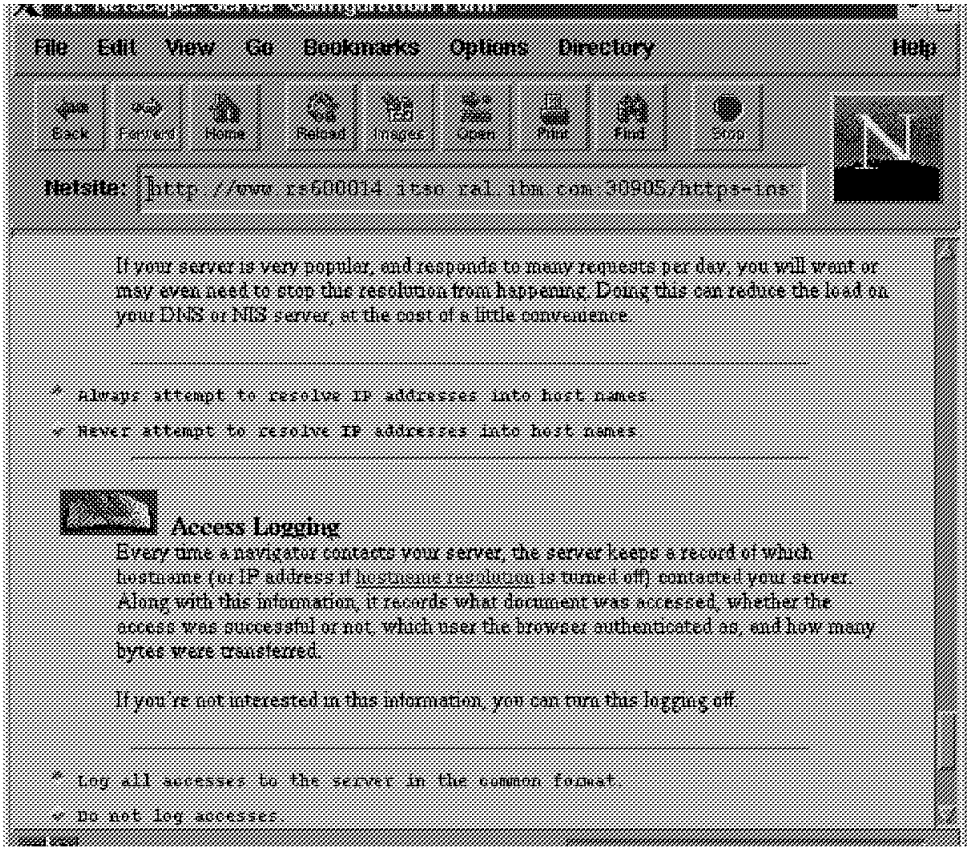


Figure 226. Web Server Installation/Configuration

When a network navigator connects to your server, the server only knows the client's IP address (for example, 204.146.46.133). The server does not know that this IP address is actually the host name www.ibm.com. For certain operations, such as access control, CGI, error reporting, and access logging, the server will resolve that IP address into a host name.

If your server is very popular and responds to many requests per day, you will want or may even need to stop this resolution from happening. Doing this can reduce the load on your DNS or NIS server at the cost of a little convenience.



Figure 227. Web Server Installation/Configuration

Every time a navigator contacts your server, the server keeps a record of which hostname (or IP address if hostname resolution is turned off) contacted your server. Along with this information, it records what document was accessed, whether the access was successful or not, which user the browser authenticated as, and how many bytes were transferred.

If you're not interested in this information, you can turn this logging off.

Click on the **Make These Changes** button in order to go to the Document Config form.

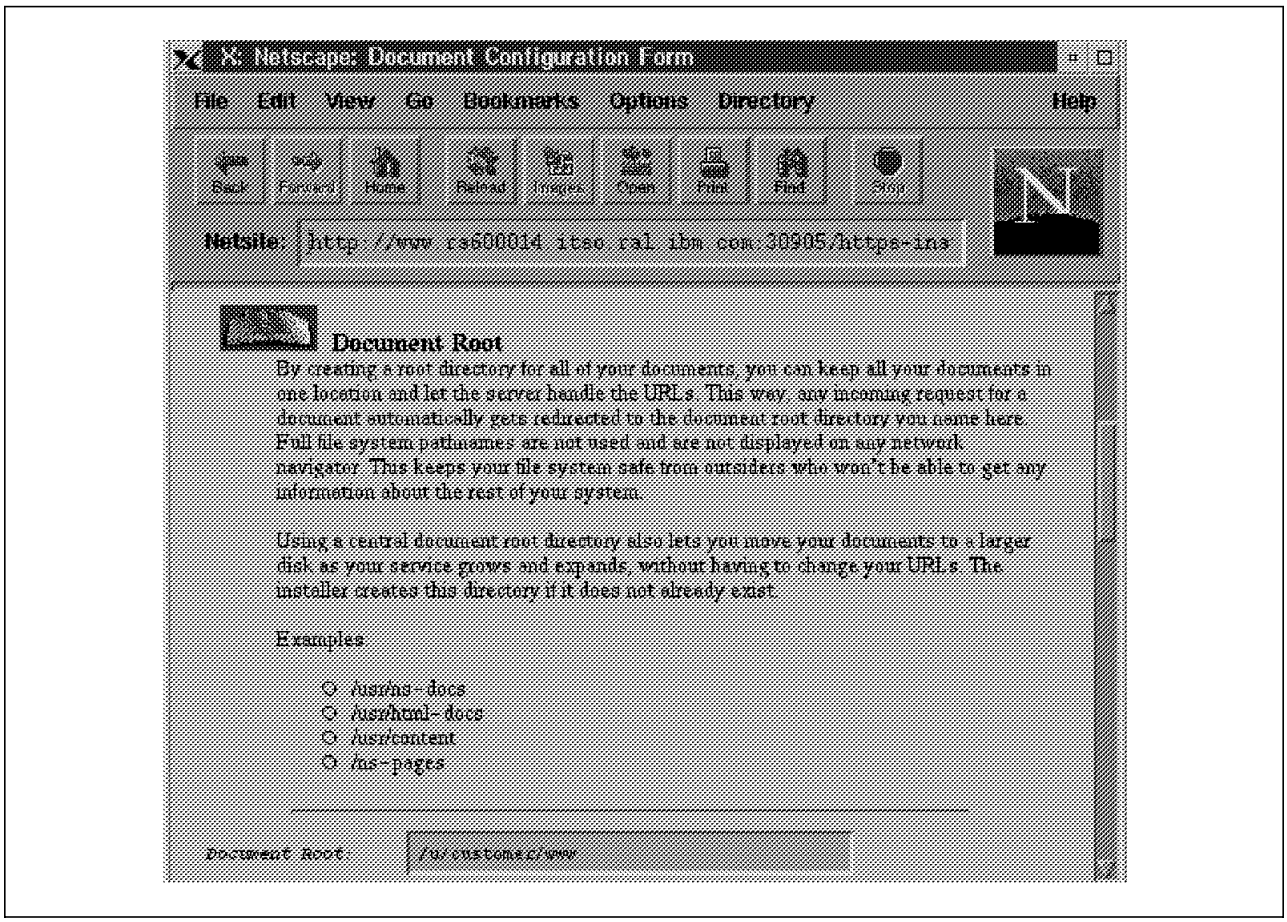


Figure 228. Web Server Installation/Configuration

By creating a root directory for all of your documents, you can keep all your documents in one location and let the server handle the URLs. This way, any incoming request for a document automatically gets redirected to the document root directory you name here. Full file system path names are not used and are not displayed on any network navigator. This keeps your file system safe from outsiders who won't be able to get any information about the rest of your system.

Using a central document root directory also lets you move your documents to a larger disk as your service grows and expands, without having to change your URLs. The installer creates this directory if it does not already exist.

Examples:

- /usr/ns-docs
- /usr/html-docs
- /usr/content
- /u/www
- /ns-pages

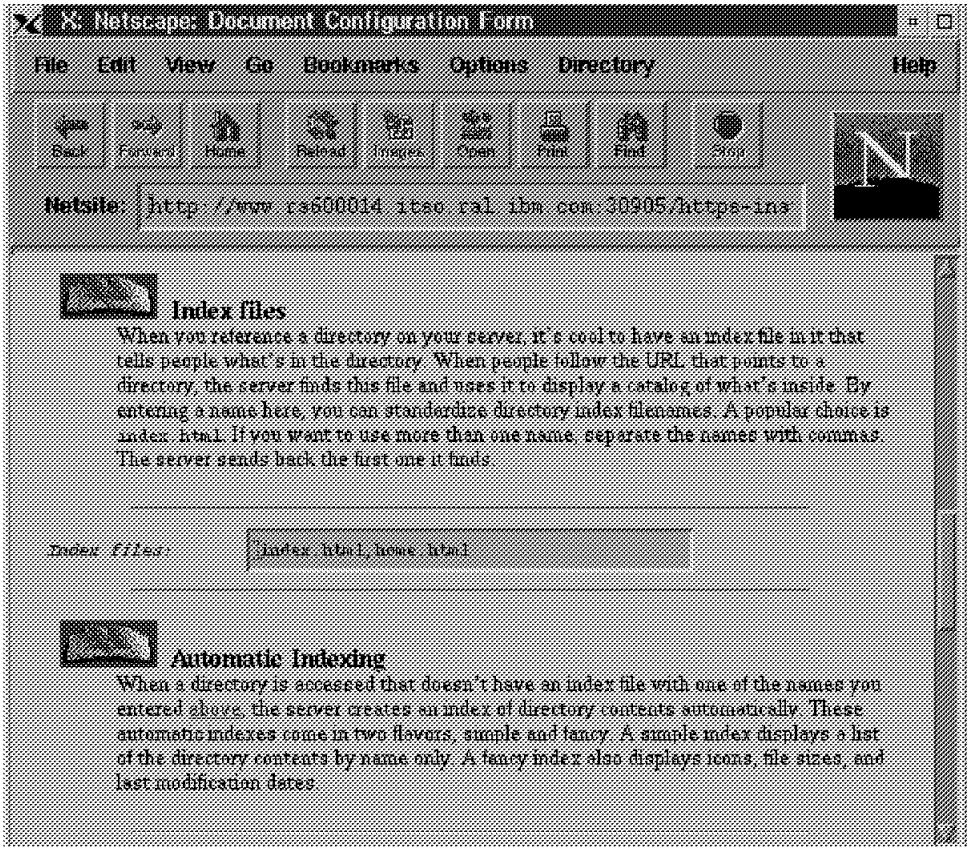


Figure 229. Web Server Installation/Configuration

When you reference a directory on your server, it's good to have an index file in it that tells people what's in the directory. When people follow the URL that points to a directory, the server finds this file and uses it to display a catalog of what's inside. By entering a name here, you can standardize directory index file names. A common choice is index.html. If you want to use more than one name, separate the names with commas. The server sends back the first one it finds.

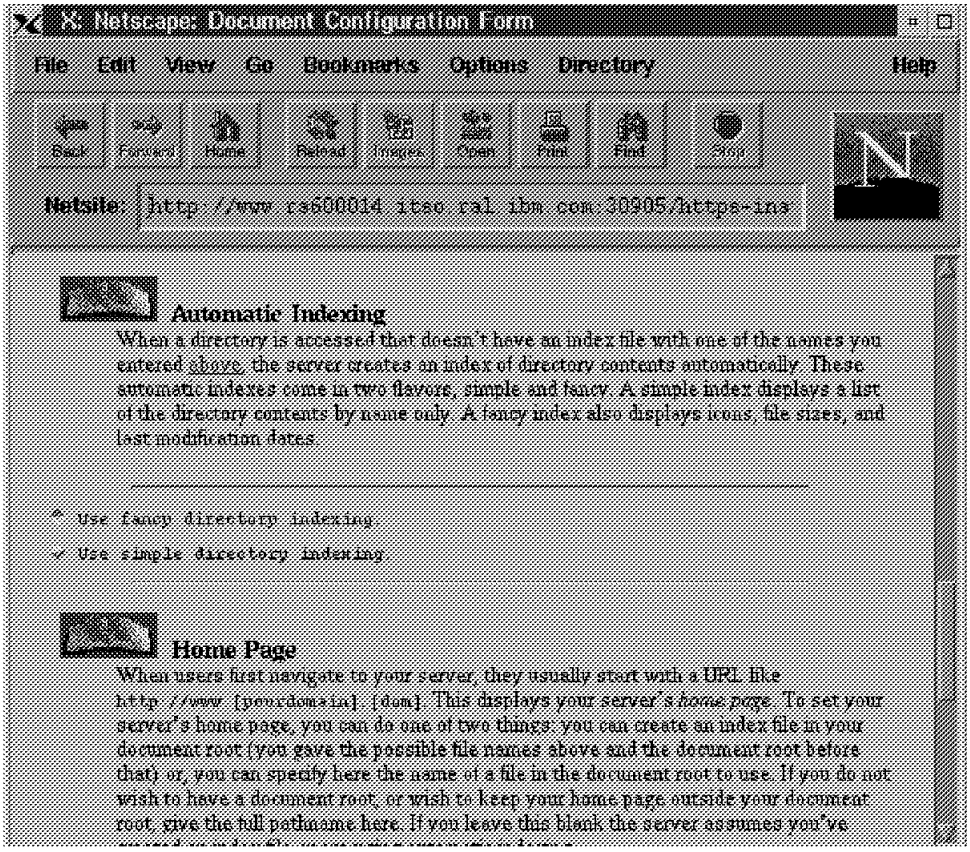


Figure 230. Web Server Installation/Configuration

When a directory is accessed that doesn't have an index file with one of the names you entered, the server creates an index of directory contents automatically. These automatic indexes come in two flavors, simple and fancy. A simple index displays a list of the directory contents by name only. A fancy index also displays icons, file sizes, and last modification dates.

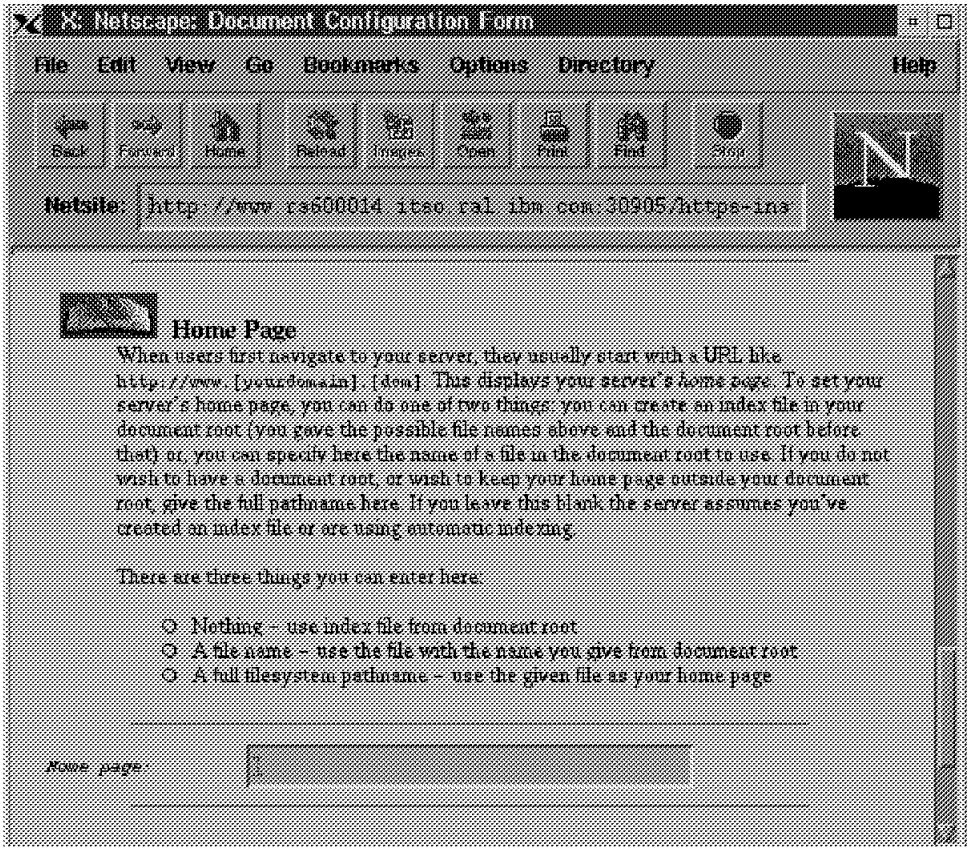


Figure 231. Web Server Installation/Configuration

When users first navigate to your server, they usually start with a URL such as http://www.yourdomain.dom. This displays your server′s home page. To set your server′s home page, you can do one of two things: you can create an index file in your document root (you gave the possible file names and the document root before that), or you can specify here the name of a file in the document root to use. If you do not wish to have a document root or wish to keep your home page outside your document root, give the full path name here. If you leave this blank, the server assumes you′ve created an index file or are using automatic indexing.

There are three things you can enter here:

- Nothing - Use index file from document root
- A file name - Use the file with the name you give from document root
- A full file system path name - Use the given file as your home page

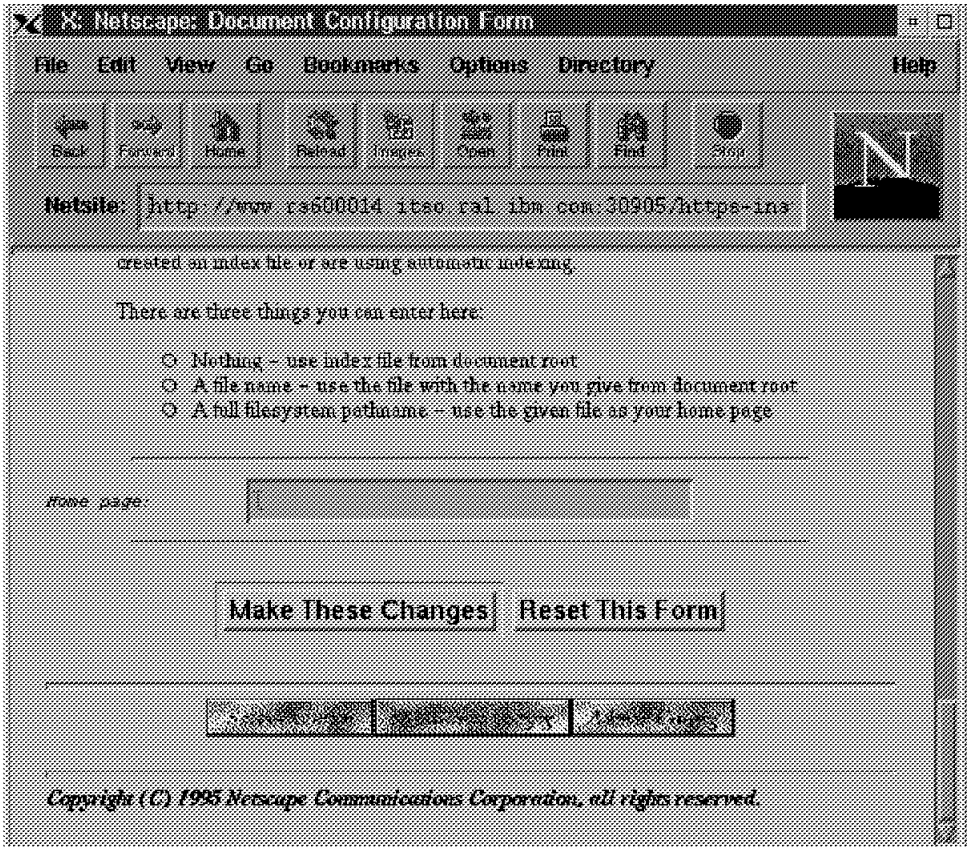


Figure 232. Web Server Installation/Configuration

After filling out the Document Config form, click on the **Make These Changes** button to enter the form information and go to the Admin Config form.

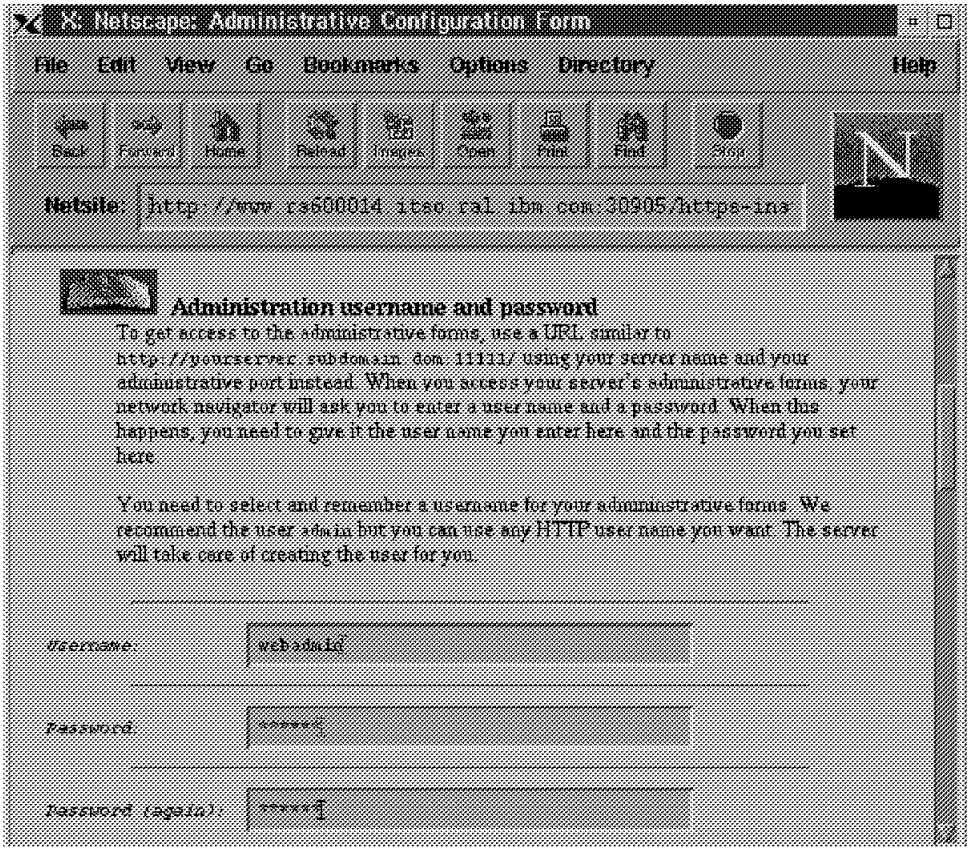


Figure 233. Web Server Installation/Configuration

To get access to the administrative forms, use a URL similar to http://yourserver.subdomain.dom:11111/ using your server name and your administrative port instead. When you access your server's administrative forms, your network navigator will ask you to enter a user name and a password. When this happens, you need to give it the user name you enter here and the password you set here.

You need to select and remember a user name for your admininstrative forms. The server will take care of creating the user for you.

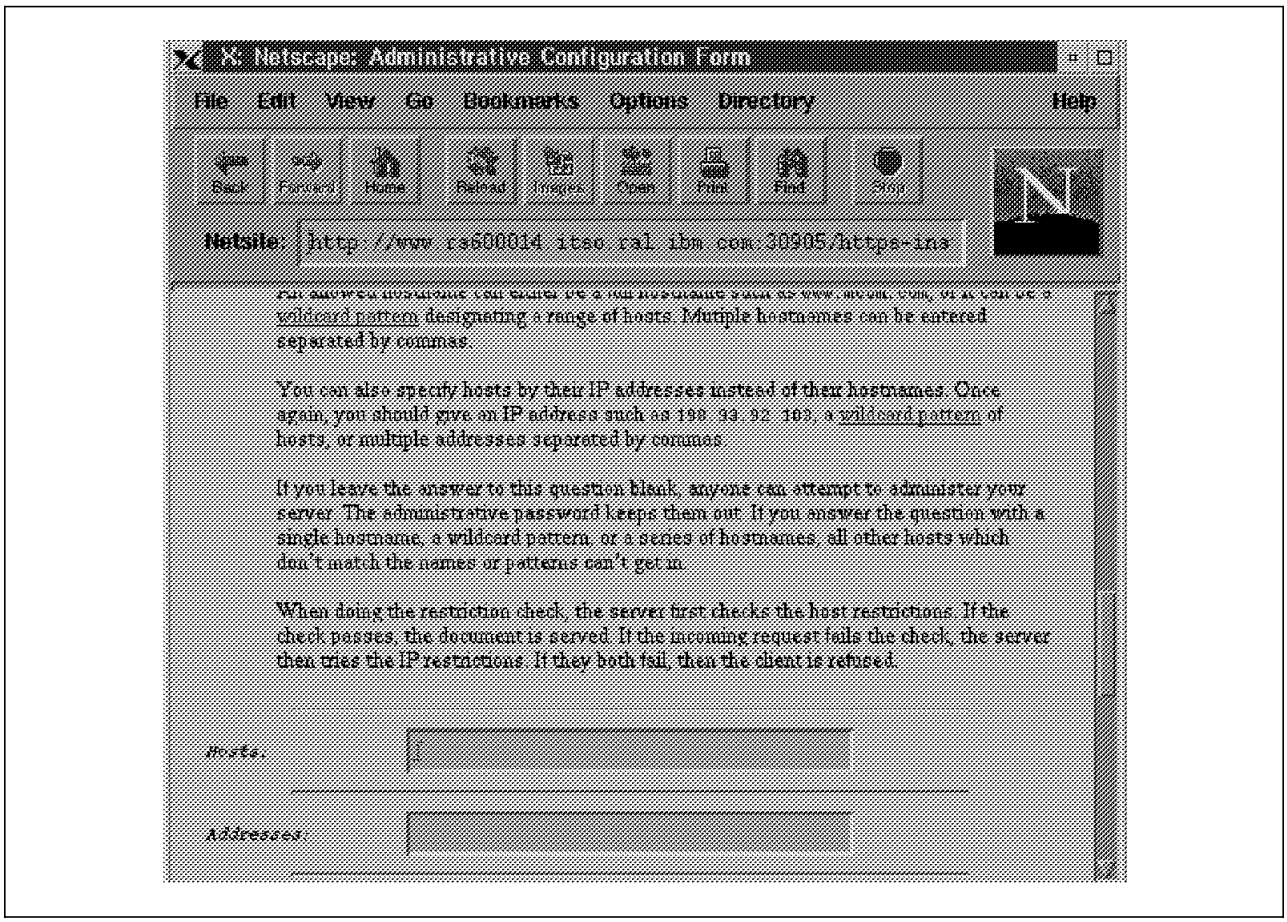


Figure 234. Web Server Installation/Configuration

Once the server is installed, you administer and manage it using your favorite forms-capable network navigator. Obviously, you don't want off-site people changing your server, and you only want authorized people to administer your server. You need to indicate here which hosts are allowed administrative access. All others will get an error if they attempt access.

If you do not trust the network between other machines and your server, you should access the administrative forms only on the server machine itself so that information never goes over the untrusted network.

An allowed host name can either be a full host name such as www.mcom.com, or it can be a wildcard pattern designating a range of hosts. Mutiple host names can be entered separated by commas.

You can also specify hosts by their IP addresses instead of their host names. Once again, you should give an IP address such as 204.146.46.133, a wildcard pattern of hosts, or multiple addresses separated by commas.

If you leave the answer to this question blank, anyone can attempt to administer your server. The administrative password keeps them out. If you answer the question with a single host name, a wildcard pattern, or a series of host names, all other hosts which don't match the names or patterns can't get in.

When doing the restriction check, the server first checks the host restrictions. If the check passes, the document is served. If the incoming request fails the check, the server then tries the IP restrictions. If they both fail, then the client is refused.

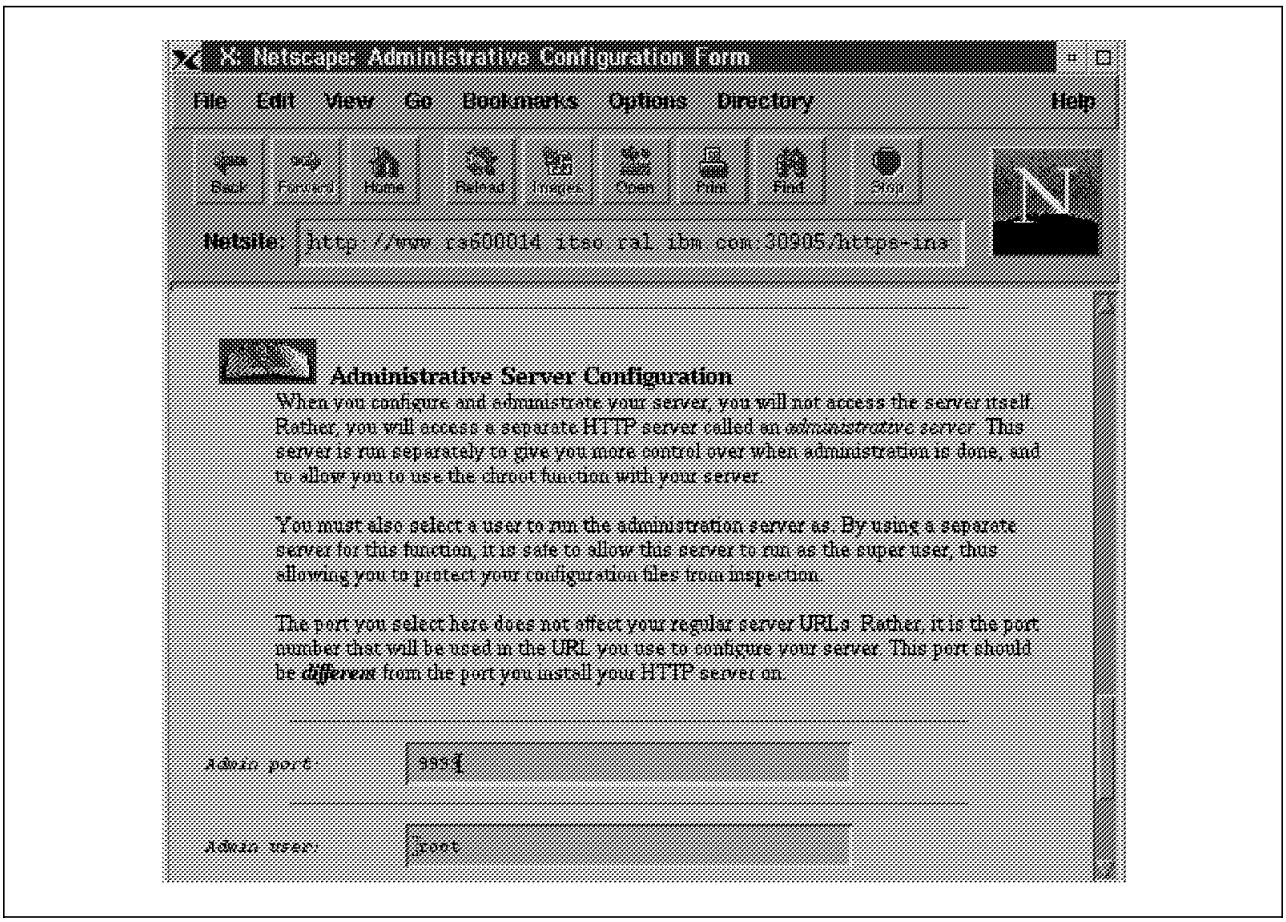


Figure 235. Web Server Installation/Configuration

When you configure and administrate your server, you will not access the server itself. Rather, you will access a separate HTTP server called an administrative server. This server is run separately to give you more control over when administration is done and to allow you to use the chroot function with your server.

You must also select a user to run the administration server. By using a separate server for this function, it is safe to allow this server to run as the super user, thus allowing you to protect your configuration files from inspection.

The port you select here does not affect your regular server URLs. Rather, it is the port number that will be used in the URL you use to configure your server. This port should be different from the port you install your HTTP server on.

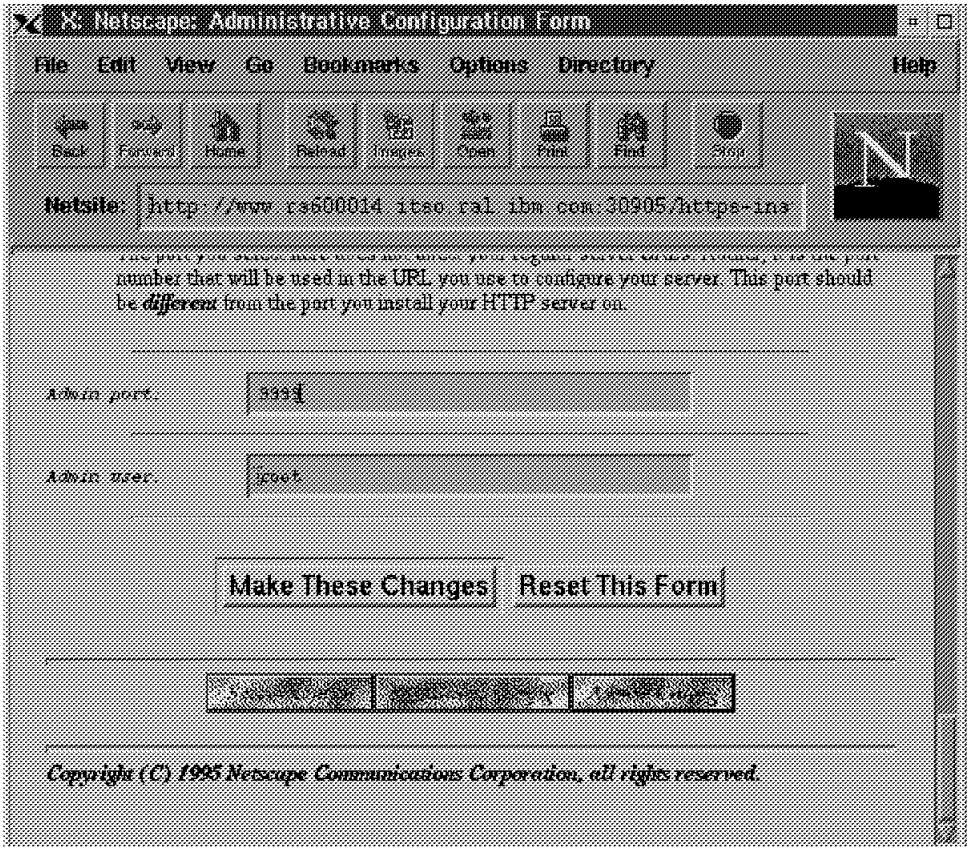


Figure 236. Web Server Installation/Configuration

After filling out this form, click on the **Make These Changes** button to enter the information and go to the installation/configuration verification process.

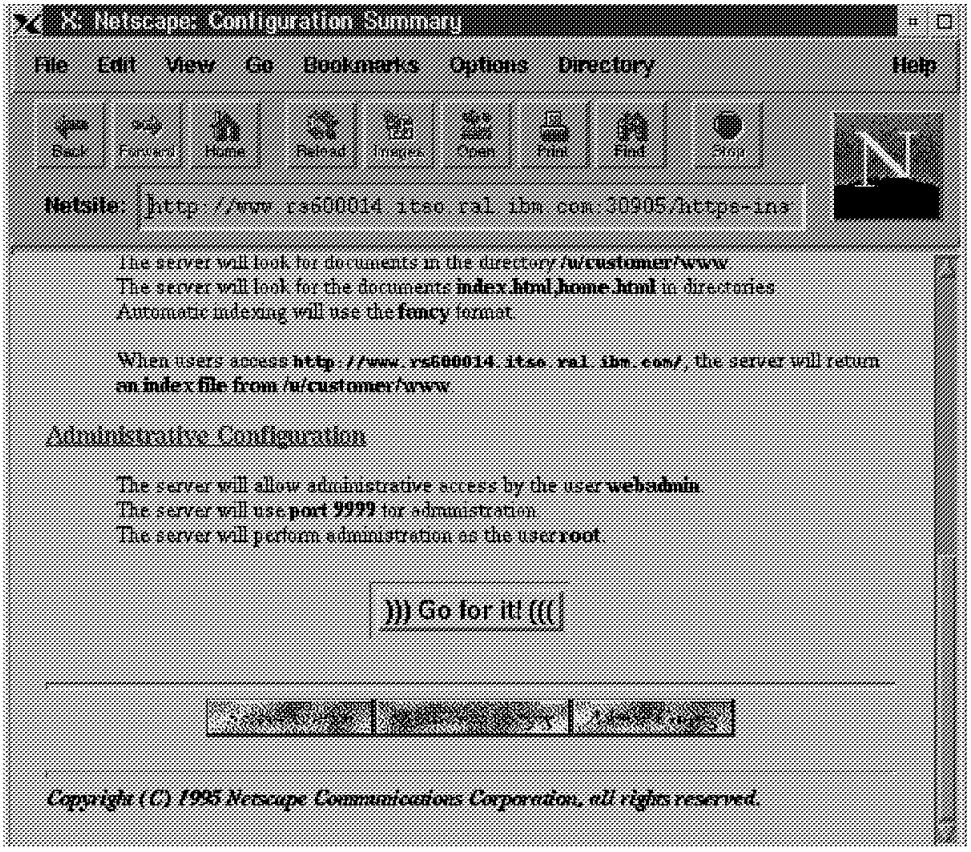


Figure 237. Web Server Installation/Configuration

After verifying the forms contents, this screen is shown with the server information summary to your last revision.

Once all the information is reviewed, click on the **Go for it!** button in order to create your customized Web server.

11.4.15 Running the Netscape Commerce Server Administration

After installation and configuration, if you need to change any item or information about your server, you must access the administration form.

To access the administration form you must enter your URL and the number of the port (Admin port; see Figure 235) you chose when you customized your server. The server will open an identification screen to enter your identification (user ID + password) that you have defined during the configuration process. Your workstation and host name or IP address defined during configuration must also be authorized to access the administrative form.

Finally, it shows a form with all the options you can use to administer your site and your customers Web servers. Each option is a specific form that you must fill out in order to request any change in the server configuration (CGI directories remapping, for example).

Using these forms you can administer the following items:

- Security Configuration. Here, you can install or renew a security certificate and enter in the advanced security items.
- Server Control. In this option you can change any information you have entered about your server configuration, such as server location, server user, number of proccesses, root directory changes, activity monitoring and server restart, startup and shutdown.
- URL Configuration. Using this item you can change your global URL configuration (server name and port and bind address) and your document configuration in order to change, for example, your server home page, default index files, etc. In this item you can also map your URLs to another server or to a local directory and view, edit and remove current URL mappings.
- User Databases. Allows you to create, edit and remove databases. Basically it is database administration.
- Access Control and Dynamic Configuration. This item must be used to restrict access by users and to configure perdirectory configuration files.
- CGI and Server Parsed HTML. This item you are going to use specifically to configure CGI applications and needs. It will be used to add, customize and activate CGI programs to your server.
- Configuration Templates. These are used to configure the templates.
- Error Handling. This allows you to view the error log and customize error responses.
- Server Administration. This item allows you to shut down or customize the administration server (admin user, port, location of the server root, and authentication user and password).
- Logging Configuration. This item is used to customize access logging.

11.4.16 Putting Web Content on the Internet

After you finish your Web server software installation, you can create a Web server for your customers whenever you want by using the same process.

Once all the customers′ environments are ready, you are supposed to contact them to inform them of their user IDs and passwords. With this information they can use FTP to access their Web servers and write their Web content in the appropriate directories.

Figure 238 shows an example of the customer default directory, with all subdirectories where CGI, image and map files must be written.

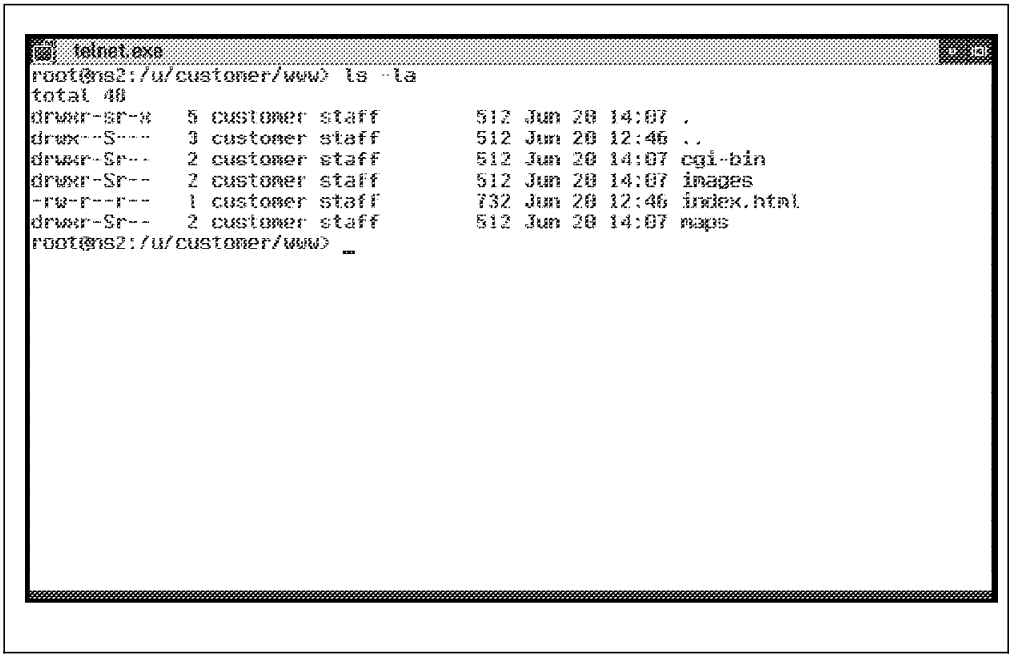


Figure 238. Example of the Customer Default Directory and Subdirectories

11.4.17 Working with CGI programs

When working with CGI programs you must create URL mappings. The next three figures show the form that must be filled out in order to add or edit URL mappings on your server. These URL mappings are used to point to directories outside of the document root.

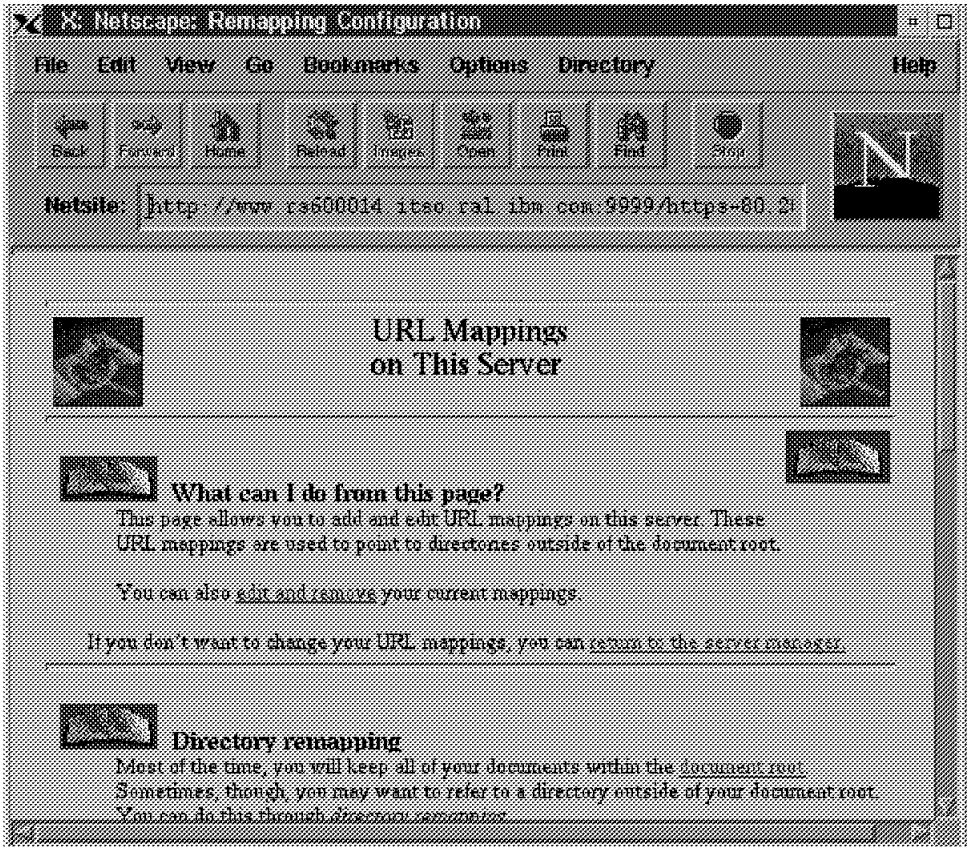


Figure 239. Web Server CGI Configuration

Most of the time, you will keep all of your documents/content within the document root, but you may want to refer to a directory outside the document root.

Use this screen to edit or remove any URL mapping or to do the directory remapping.

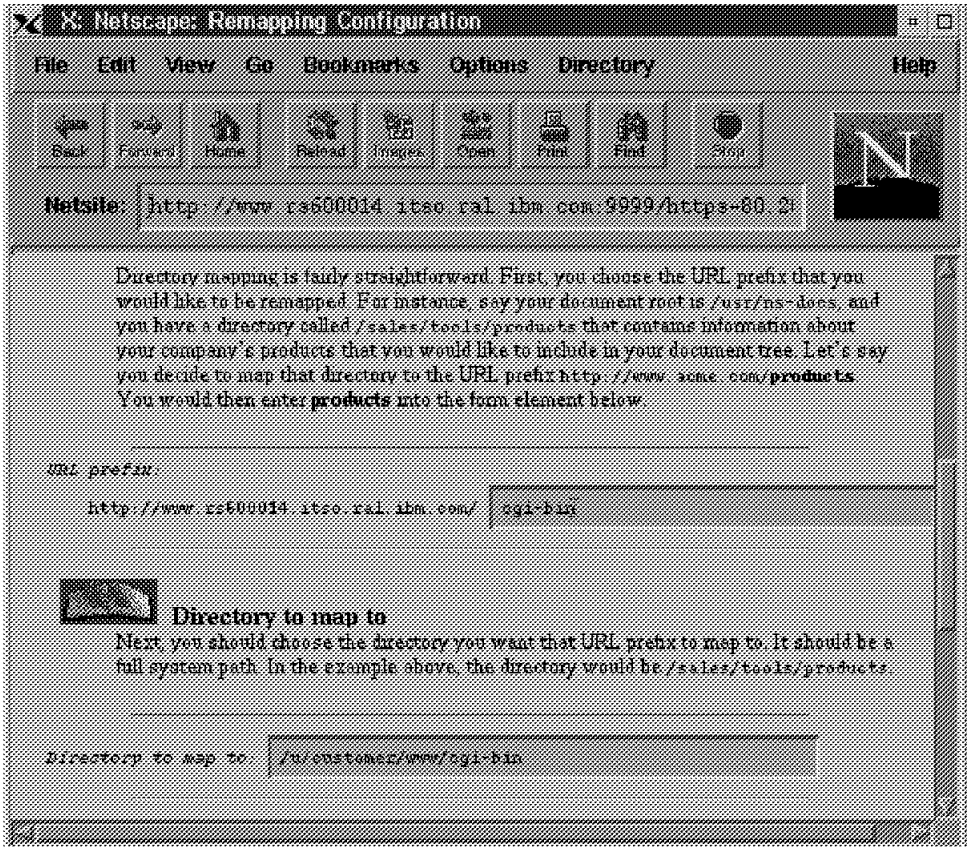


Figure 240. Web Server CGI Configuration

In this screen you choose the URL prefix that you would like to be remapped. For example, say your document root is /usr/ns-docs, and you have a directory called /sales/tools/products that contains information about your company's products that you would like to include in your document tree. Let's say you decide to map that directory to the URL prefix http://www.acme.com/products. You would then enter your URL prefix and the directory path to map to in this screen's fields.

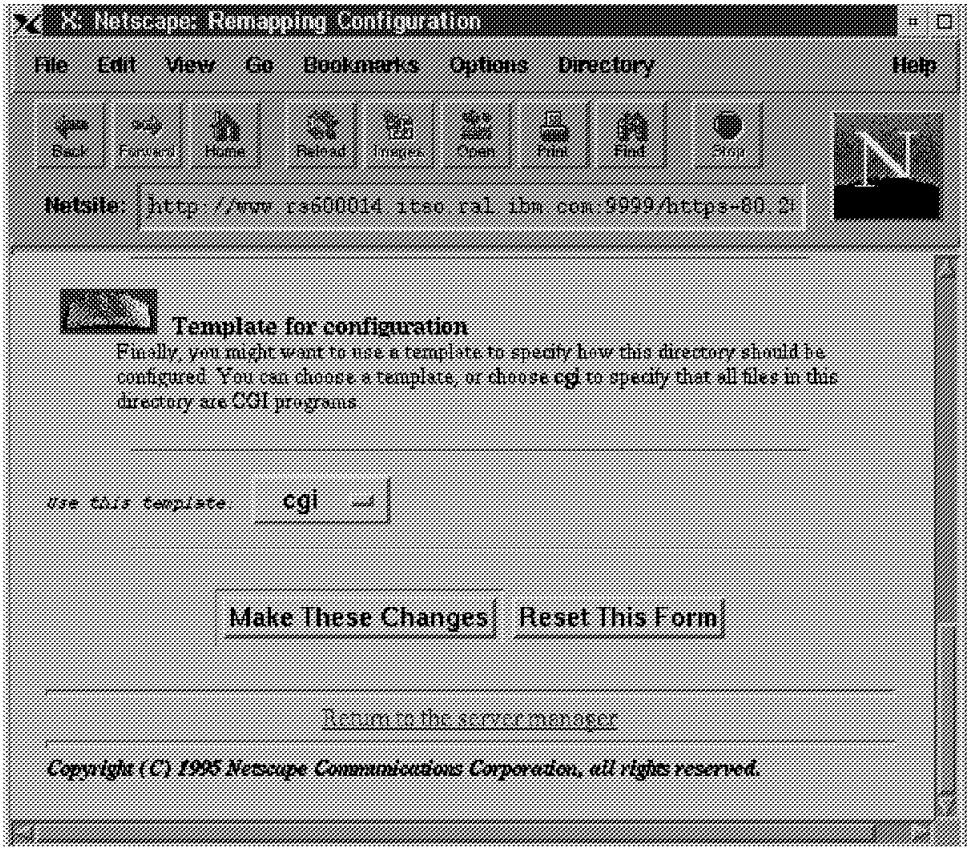


Figure 241. Web Server CGI Configuration

The third screen is used to make the template for configuration, because you might want to use a template to specify how this directory should be configured. You can choose a template or a CGI to specify that all files in this directory are CGI programs.

Then you must click on the **Make These Changes** button to finish the CGI directory creation process.

11.4.18 Developing an FTP Site

Lots of systems connected to the Internet have file libraries or archives accessible to the public. Most of these consist of free or low-cost shareware programs for virtually every computer.

The File Transfer Protocol (FTP) is different from Telnet, because FTP will let you bring the information down to your computer. You can then see it or use it whenever you want, regardless of the remote site availability or the speed of the communication lines.

FTP is widely used for transferring files between computers on a TCP/IP-based network such as the Internet. Its login will also ask you for a user ID and a password. In most cases, the user ID can be *anonymous* and the password can be your Internet e-mail address.

For security reasons, anonymous FTP is implemented on Internet servers as a way to access publicly available software and files. When developing an

FTP site it is interesting to create a customers′ anonymous FTP area where they can put their files for public distribution. This section explains how the anonymous FTP area works and how to set it up in our sample AIX (IBM UNIX) environment.

11.4.18.1 Understanding the Anonymous FTP Area

First, the FTP daemon must be made available to the AIX operating system and all the directory structures must be created under a restricted sample directory. One of these directories must be called pub and under it all the customers′ downloadable files must be kept.

The anonymous FTP area is a restricted area for the anonymous FTP user. When someone FTPs into the server as anonymous, they are put in a directory environment (for example, chroot directory). This chroot environment is a restricted environment designed to allow only user access to the directories below the chroot sample directory.

Figure 242 shows an example of a chroot environment with the chroot sample directory being /anonftp.

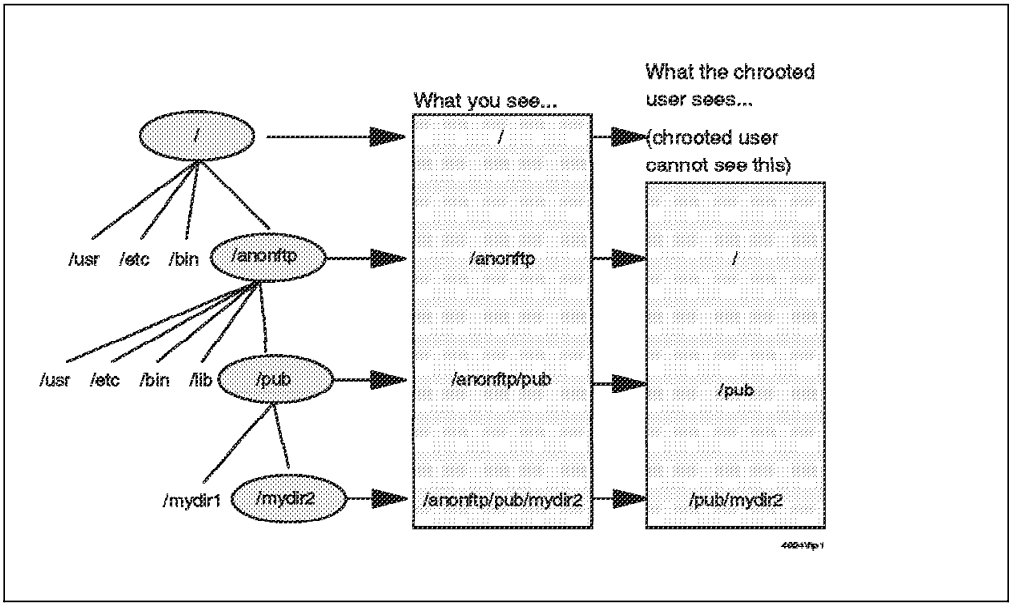


Figure 242. chroot Environment Example. The chroot sample directory is /anonftp.

Notice there are two levels of system-type directories (/usr, /etc, and /bin) shown in Figure 242. It is very important to understand the difference between the two levels of system-type directories. In order to help you understand the function of each level of the directory structure, the functional differences are explained in the following text and graphically represented in Figure 243 on page 514.

- Your system-level directories. These directories determine your customers system configuration. All the binaries that they are going to use are in the /bin directory, their configuration files are in /etc, and so on. The customers must not have write access to this level of directories.
- The anonymous FTP area system-level directories. These directories are used only by the anonymous FTP users. For example, the executables that are in /anonftp/bin are the only executables available to anonymous FTP users. Because the customers are provided with a minimally

configured setup, the only executable they will find in the /anonftp/bin directory is ls (UNIX command to list files). If they want any other executables to be available to anonymous FTP users, they have to copy them from /bin to /anonftp/bin. This is the same for the /usr, /etc, and /lib directories. They must have complete write access to this level of directories.

- The anonymous FTP area public directories. This directory level is not for system-level directories; it serves as the directory-level where the customers can put publicly downloadable files. They can create directories and files on this level. By default, no directories are created on this level.

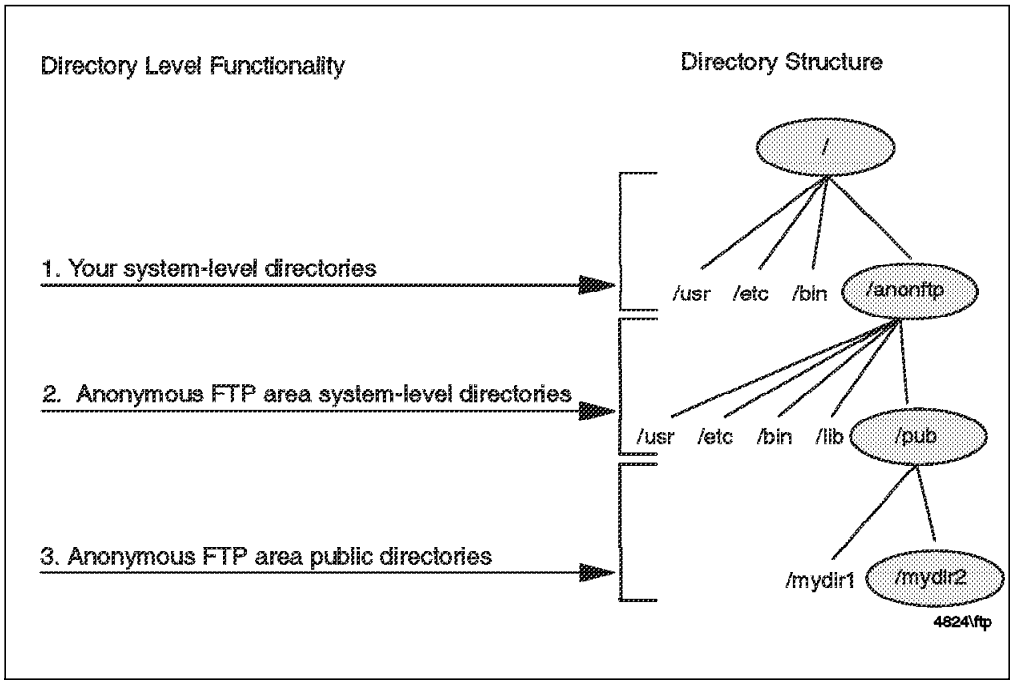


Figure 243. Function of Each Level of the Directory Structure of the Anonymous FTP Area

11.4.18.2 Configuring Your Anonymous FTP Area

Your customers′ anonymous FTP area is preconfigured. If they have files available to download, this preconfiguration allows users to FTP into their servers and download files. Further configuration options include:

- Having other commands such as tar and uncompress available to the anonymous FTP users.
- Tailoring the FTP password and group files to customize your customers′ setup.
- Adding an incoming area so anonymous FTP users can upload files.

To make other commands available to anonymous FTP users, copy the binary from your customers′ /bin directory (their system-level directory) to the /anonftp/bin directory (anonymous FTP area system-level directory). Make sure that the binary has executable permissions (---x--x--x).

Your customers may also want to tailor their FTP passwords and FTP group files to customize their setups. The entries in the FTP password file (/anonftp/etc/passwd) and the FTP group file (/anonftp/etc/group) do not

affect who has access to the servers. The function of these files is to show a logged-in anonymous FTP user file and directory ownership names instead of the user's ID numbers.

Figure 244 shows what the anonymous FTP user sees when the password and group files are implemented. The format is easier to read and shows which company is providing the service.

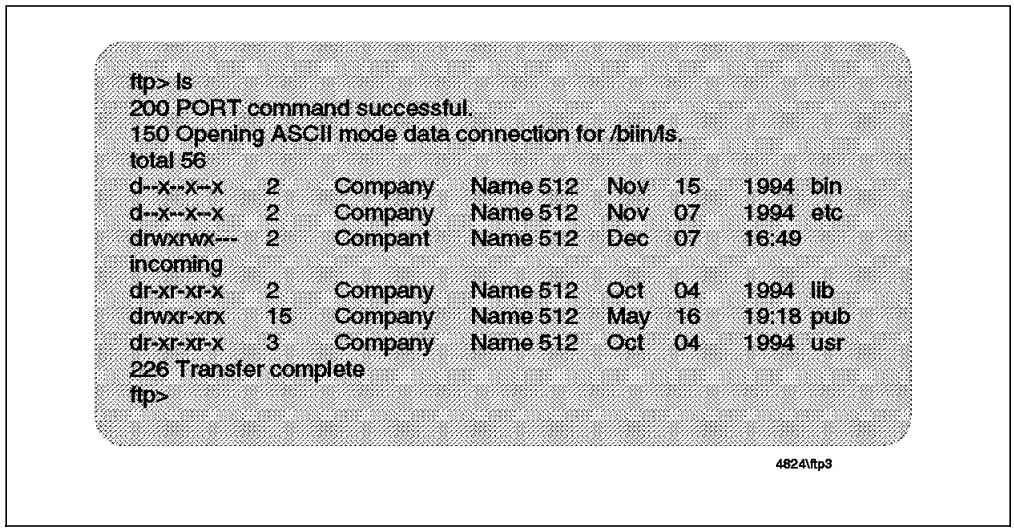


Figure 244. Entries for the Password and Group Files Implemented

11.4.18.3 Creating an Incoming Area Where Anonymous Users can Upload Files

Your customers may want anonymous users to be able to upload files to their servers. This can be done, but it can be risky and requires some administrative overhead on your part.

Having an upload area can be risky because you have little control over the content that someone uploads. They could fill all your available space, which would deny logging capability and other services relying on space.

The following are recommendations for setting up an incoming area.

1. Create a directory for example:

   ```
   mkdir /anonftp/incoming
   ```

2. Protect the top-level incoming directory, giving only execute permission to the anonymous user for example:

   ```
   chmod 751 /anonftp/incoming
   ```

3. Create subdirectories in the /anonftp/incoming directory, using names known only to people who are allowed to upload for example:

```
mkdir /anonftp/incoming/NnSf4j

chmod 753 /anonftp/incoming/NnSf4j
```

This way only people who know the specific directory name can cd (change directory command) into it. This directory is the only place with upload permissions.

11.4.19 Getting Reports from Content Services

One of the softwares we said you must have installed in your server is the Web server report software; that is, a software to get statistic reports about your Web server behavior.

If you need this kind of report to manage your Web site, your customers also need the reports to evaluate the availability of their services, which home pages are being accessed more, which ones are being accessed less, what kind of things are being transferred from their pages, from where these accesses are being made, and the amount of accesses in each hour of the day, day of the week, week of the month and month of the year.

These kind of reports usually bring the number of hits of each home page in your Web server. They run using your Web server log information, which keeps all the information on the number of requests addressed to your server along with other information, such as the time of day the requests were made.

You can use in your Web server any software that can give to you and your customers the information that both of you need. Some of these softwares are free, while others are sharewares or even more expensive but are more user friendly and allow graphic configuration.

11.4.19.1 Getting Web Server Reports

You must inform your customers of when they can request their reports (usually weekly or monthly). Depending on the software that is being used, you must inform your customers by e-mail or give them a user ID and password to access a statistic report home page generated by your software. Reports must always be automatically sent to your customers.

11.4.19.2 Producing Their Own Reports

If your customers have special needs that are not met by the standard reports, they must be able to use their log information and generate their reports using their own Web server report software.

In order to do this your customers must be informed of where to locate the log files. The default option is to create a log directory under your www (the customer's default Web server directory). For example:

```
/www/log
```

11.4.19.3 Using Statistics Report Options

The getstats program is freely available on the Internet. You and your customers can learn more about it and other report-generating software at the URL http://www.eit.com/software/getstats/getstats.html.

A graphic option can be the WebTrends software that generates a report home page. Further information about this software can be found at the URL http://www.webtrends.com.

The following tables are examples of the getstats program reports.

Table 38. Header Information

Option	Default	Example
-c	yes	HTTP Server General Statistics Server: http://www.hostname.com/ (NCSA Common) Local date: Thu Oct 12 14:54.12 PM EDT 1995 Covers: 09/27/95 to 10/01/95 (5 days) All dates are in local time. Requests last 7 dyas: 0 New unique hosts last 7 days: 0 Total unique hosts: 9649 Number of HTML requests: 53172 Number of script requests: 13931 Number of non-HTML requests: 185567 Number of malformed requests (all dates): 39599 Total number of all requests/errors: 292269 Average requests/hour: 2828.1, requests/day: 67875.1 Average bytes/hour; 14110626, bytes/day: 338655025

Table 39. Daily Summary

Option	Default	Example
-ds	no	HTTP Server Daily Summary Covers: 09/27/95 to 10/02/95 (6 days) All dates are in local time. Each mark (#) represents 1600 requests ---------- Mon: 29474 : # Tue: 38840 : ## Wed: 60416 : #### Thu: 79124 : ###### Fri: 74904 : ###### Sat: 39750 : ## Sun: 38402 : ##

Table 40. Daily Report

Option	Default	Example
-d	yes	HTTP Server Daily Statistics Covers: 09/27/95 to 10/01/95 (5 days) All dates are in local time. Each mark (#) represents 1600 requests ---------- 9/27/95 (Wed) : 20416 : ############# 9/28/95 (Thu) : 79124 :#########################\|\|## 9/29/95 (Fri) : 74904 : #########################\|\|## 9/30/95 (Sat) : 39750 : ######################## 9/01/95 (Sun) : 38402 : ########################

Table 41. Hourly Report

Option	Default	Example
-h	yes	HTTP Server Hourly Statistics Covers: 09/27/95 to 10/02/95 (6 days) All dates are in local time. Each mark (#) represents 64 requests ---------- 9/27/95 (Wed) 4:00pm 29474 : ########## 5:00pm 29474 : ########## 6:00pm 29474 : ########## 7:00pm 29474 : ########## 8:00pm 29474 : ########## 9:00pm 29474 : ########## 10:00pm 29474 : ########## 11:00pm 29474 : ########## total: 20416 9/28/95 (Thu) midnite 29474 : ########## 1:00am 29474 : ########## 2:00am 29474 : ########## 3:00am 29474 : ########## 4:00am 29474 : ########## 5:00am 29474 : ########## 6:00am 29474 : ########## 7:00am 29474 : ########## 8:00am 29474 : ########## 9:00am 29474 : ########## 10:00am 29474 : ##########

Table 42. Request Report Sorted by Date

Option	Default	Example
-rd	no	HTTP Server Request Statistics Covers: 09/27/95 to 10/02/95 (6 days) All dates are in local time. Sorted by last access date, 2548 unique requests # of requests : Last Access (M/D/Y) : Bytes/File : Request ---------- 8149 : 10/02/95 : 4727016 / 581 : /images/picture1.gif 455 : 10/02/95 : 2037945 / 4479 : /dir/file1.html 5243 : 10/02/95 : 29711249 / 5717 : /images/picture1.gif 1158 : 10/02/95 : 3682440 / 581 : /images/picture1.gif 13578 : 10/02/95 : 4727016 / 581 : /images/picture1.gif 1253 : 10/02/95 : 4727016 / 581 : /images/picture1.gif 1292 : 10/02/95 : 4727016 / 581 : /images/picture1.gif 4156 : 10/02/95 : 4727016 / 581 : /images/picture1.gif 11208 : 10/02/95 : 4727016 / 581 : /images/picture1.gif

11.4.20 Network Monitoring and Management

This section covers some network monitoring and management procedures that need to be implemented.

11.4.20.1 Creating a Management Process

When talking about management we need to imagine a big process where you describe all your needs in order to have tools to actually make your Web site administration. You need to be aware of the Web server management software you are going to use, what will be your network management, what kind of services monitoring process you are going to perform, how to create a help desk service and the management and monitoring reports you need to receive daily to analyze your services′ availability.

11.4.20.2 Web Server Management Software

What is the purpose of a virtual marketplace if the doors are shut and locked when the shoppers want to shop or when you need important information from a Web server? The Web server management software is designed to track the availability and performance of critical Internet server devices, avoiding your services′ interruption.

To use a good web server management software, you need to know what it is supposed to do. In order to help you, we have listed some of the features this software must have. It must:

- Detect all devices in a TCP/IP network.
- Monitor all devices in a TCP/IP network.
- Test applications of the Internet Web servers.
- Test the server functionality of the Internet Web servers.
- Inform availability and performance.
- Servers tested must include at least:
 - WWW
 - FTP
 - Mail (SMTP)
 - News
 - Gopher
 - Archie
 - WAIS

An example of this kind of software is the Caravelle′s WebWATCHER. Further information about the features of this software can be found at the URL http://www.caravelle.com/noframes/web.htm#WebWATCHERTM.

11.4.21 Network Management

Not only the TCP/IP devices and Internet servers need to be monitored, but also all the network(s) where they are plugged need to be monitored.

Before deciding which software or tools you need to buy you must understand your environment. Depending on your environment you need to use different network management solutions. For further information about this, refer to the Chapter 9, “Network Management” on page 385.

11.4.21.1 Monitoring Process

In order to monitor your Web site servers, software and tools are important but you need to have a help desk team. You are supposed to create a group able to analyze and make a quick decision in case of any problem.

If the Web server management software informs you that something is not working well, this team is supposed to solve the problem, even if they need to contact any other technical specialist. This is necessary to guarantee your service availability and reliability.

As we mentioned in the begining of this document, you need to offer this service 24 hours a day, 7 days a week in order to meet your customers′ needs, especially when your customer has a commercial site that sells goods. You need to guarantee this service because the users can shop at any time. This is the real Internet user new profile.

11.4.21.2 Management and Monitoring Reports

All the software and tools you are going to use must return to you eventually (we recommend daily) reports about your services′ availability and problems. This will help you in your troubleshooting process and to analyze your services′ growing needs.

11.4.22 Operational Issues

Besides the management and monitoring process you must have an operational team to handle customers′ doubts or help them to request some administrative or technical services.

This service must also be available 24 hours a day, 7 days a week and must be done for another team or even the help desk you created to monitor your services. But in this case, an operational process is necessary to avoid any misunderstanding.

11.4.23 Security

The only way to guarantee Web server security in your environment is limiting access to your Web servers. This topic serves as a reminder to setting up access control to your Web server, so you can control access to the production areas of your servers. The following are three ways to control access to your Web content:

- Limiting access by the customer′s Internet domain name.
- Limiting access by user.
- A combination of both.

11.4.23.1 Limiting Access by the Customer′s Internet Domain Name

With this option your customers can allow or deny access to users by their Internet domain name. For example, they can allow users from their company′s domain to access their areas through the Web. They would do this to keep anyone outside of their organization from seeing their Web pages before they were ready for general use. They can limit access by domain name in two ways: allow access from a domain or deny access from a domain.

Two issues are relevant in deciding how to limit access by domain name. The first is to decide which Web pages your customers want to control access to. The second is to decide which domains they want to allow and which they want to deny access to their Web contents. Limiting access to certain Web pages can only be done on a directory-tree basis. In other words, they can limit access to Web pages in a certain directory and all of its subdirectories.

11.4.23.2 Limiting Access to Only Certain Users

Your customers may also be able to allow access to only certain users. To do this, a password file that contains the users they want to access their Web content must be created.

Notice the difference between user access to customers′ Web content and user access to the server. Server access and Web content access are completely independent, which means that a user having access to one does not necessarily have access to the other. Server access (user names and passwords) is controlled by the server administrator, but Web content access must be controlled by your customers. Because server access and Web content access are independent, a user who has access to the server and the Web content can have different user names and passwords for each.

11.4.23.3 System Administration Controls

The security bit permissions (talking about UNIX systems) must be created to avoid unauthorized access to your Web server by any customer or even a hacker or cracker.

Any security problem well-known by the system administrators must be eliminated in order to avoid operating system environment violation attempts. You must guarantee that your customers only have access to their own directories and subdirectories for controlling access to their Web pages.

Another thing to be aware of is the passwords creation. You must create rules to avoid the occurance of trivial passwords and to force the passwords to be changed. The process that informs your customers of a new password (for the first time or in case of loss) must also be controlled to avoid any security disclosure.

For further information about security in the Internet environment, refer to the Chapter 8, “Security on the Internet” on page 339.

Chapter 12. Networked Applications

In this chapter we describe networked applications available in the Internet like IBM infoMarket and IBM infoSage. We also describe the Net.Commerce product that allows you to build your own networked application in the Internet.

12.1 IBM infoMarket

The IBM infoMarket service is part of the IBM network computing strategy. IBM infoMarket combines a sophisticated search engine with an unmatched body of content, secure container technology and IBM′s expertise in transaction and billing management, to provide a framework for secure electronic commerce. The Internet is a natural conduit for the transmission of information, but publishers previously have been afraid to put their valuable content on the ′Net for fear that it may be tampered with or that they will not receive compensation for their works. The IBM infoMarket service provides the technology to address those needs.

With a growing number of content and technology providers now aligned with the service, including Reuters New Media, Jupiter Communications and CMP Publications, IBM infoMarket is a leading provider of content on the Internet. Users benefit from IBM infoMarket′s ability to search through commercial content some of which was previously unavailable on the World Wide Web.

Once targeted information is found, commercial content will be delivered in Cryptolope containers, accompanied by a content abstract that provides users with the essential product knowledge they need to make a buying decision. The content abstract may also include the content source, summary, author, last update, size, and price, as well as any unique terms of sale. Once the user has decided to open the contents of a Cryptolope container, a transparent digital key is issued unlocking the material contained within.

To view a free document, the user clicks on the article and the information appears on the desktop. To view priced content, the user agrees to any unique terms of the Cryptolope container that are prepared by the rights holder and contained in the content abstract.

Users of IBM infoMarket will be able to retrieve and download Cryptolope containers from the infoMarket service Web site using an IBM infoMarket helper application, initially for Netscape 1.1 or greater, and Windows 3.1 or Windows 3.1.1. Windows 95, OS/2 and Mac/OS versions are planned.

© Copyright IBM Corp. 1996

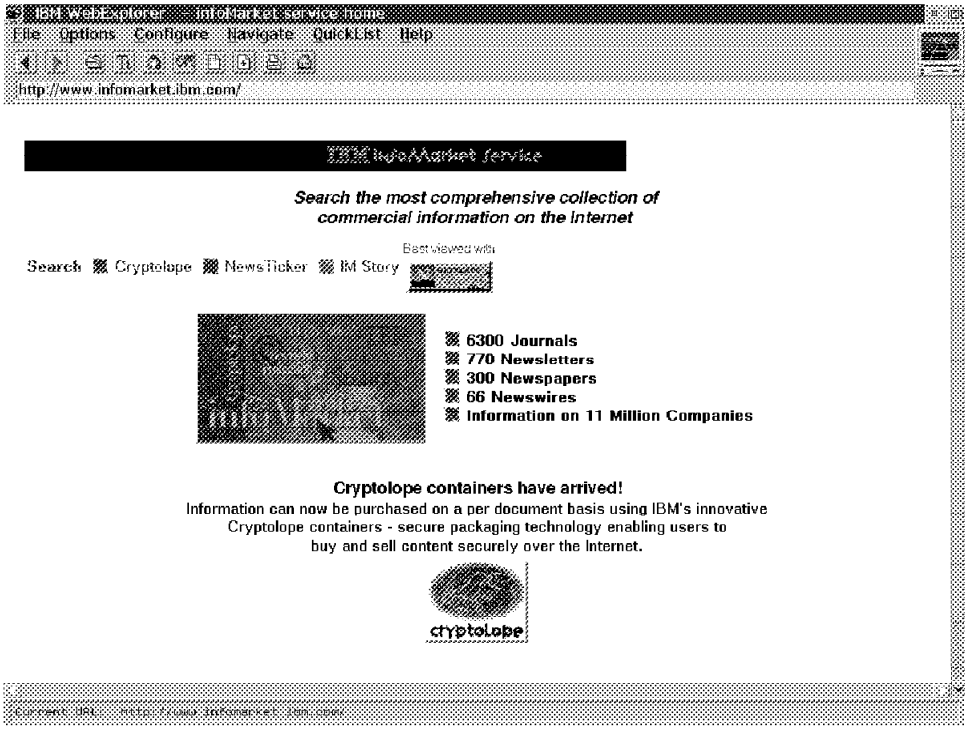


Figure 245. IBM infoMarket Home Page

12.1.1 Wide Area Search of Distributed Data

The search feature of the IBM infoMarket service can take a single search request and simultaneously search multiple, disparate, and distributed information sources returning relevance-ranked results. IBM infoMarket's powerful search capabilities allow users to simultaneously investigate the broad scope of the Internet as well as authorized, private commercial content. Keyword and Boolean search requests allow for more complex information queries using OR, AND or NEAR connectors. These search capabilities allow for manual selection of sources that will best contain the desired information to meet users needs.

IBM infoMarket also provides a feature which enables users to search by source. Users can pick an individual source, all or a group of sources in a category, or select all sources. Additionally, search results can be received in brief or with more details. Detailed results provide the user with the content abstract, which may include source, author, excerpt, date, size, and pricing information.

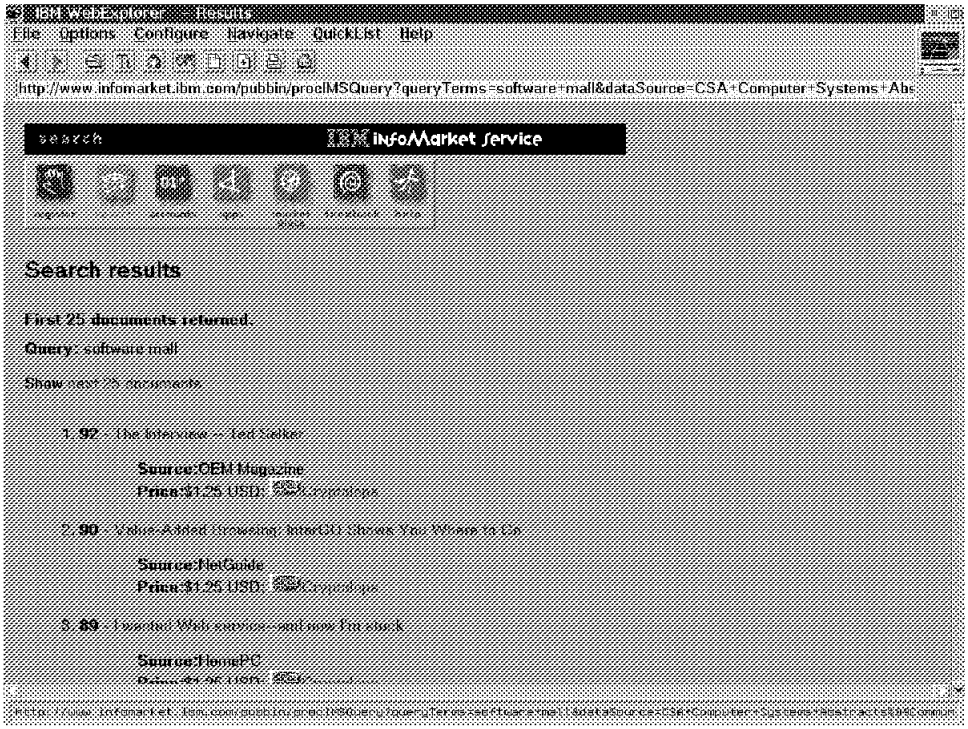


Figure 246. IBM infoMarket Search Result

New content and technology providers for IBM infoMarket include the following:

- CMP Publications provides publishing, marketing, and information services to the high-technology market
- Excalibur Technologies Corporation provides the first set of tools for creating retrieval solutions for text, images, and other forms of digital data
- International Services, Inc. provides international trade leads, language translation services, and international credit reports
- Jupiter Communications, provides LLC research, consulting, and publishing information emerging consumer online and interactive technologies.
- M.A.I.D established Profound, Inc., the New York-based subsidiary of Market Analysis and Information Database, Inc.
- Market Guide provides high quality, fundamental information on over 8,000 publicly traded companies to the professional brokerage, institutional research and individual investor marketplaces.
- Online Inc. publishes how-to magazines and books aimed at users of online databases, CD-ROMs, multimedia, and the Internet.
- Thunderstone Software provides intelligent concept searching and retrieval technologies.
- Vickers Stock Research Corporation, Inc. provides information on exactly what stocks insiders are buying and selling.

12.1.2 Cryptolope

The objective of the IBM infoMarket rights management system is to maximize revenue from the controlled use of information assets by deterring economic loss because of piracy and other unauthorized use, while enabling superdistribution. It is very important to recognize that the objective is not to prevent piracy altogether. Indeed, current technology cannot prevent piracy entirely, as recent hacks into ″secure″ systems have demonstrated. The usage of information assets without authorization or payment is a continuing threat. To minimize the economic impact of this threat, the IBM infoMarket rights management system seeks to decrease the number of security breaches, the scope of damage of each breach, and its spread. By minimizing the exposure in these three areas, the IBM infoMarket rights management system assures information providers that they will receive maximum revenue for their content, given reasonable costs to prevent unauthorized use.

12.1.3 Cryptolope Container

The mechanism that the IBM infoMarket service will implement to control and monitor the use of information is based on a secure container architecture for packaging and distributing information content and properties. We call this container a cryptographic envelope, or Cryptolope container. A Cryptolope container holds an encrypted version of a document (may contain many data formats such as ASCII text, HTML, image, and so forth) as well as rules for determining permissions specified by the content provider.

A Cryptolope container also holds control information that describes the document contents such as an abstract, price, and restrictions or terms and conditions on the use of the content. This control information is available without decrypting the actual document contents.

The data in the Cryptolope container is cryptographically signed to prevent undetected alteration of it.

12.1.4 Key IBM infoMarket Rights Management Directions

The IBM infoMarket rights management system design is driven by the following principles:

- Content providers specify rules for controlling access permissions, costs, and document restrictions. Facilities will be provided to help content providers generate these rules and permissions (that is, through a software application). These rules are attached to the content and remain with it within the Cryptolope container as the Cryptolope contents are distributed.
- Cryptolope allows for flexible specification of these rules and how the rules can be evaluated.
- At the time of content/service access, the user credentials, the requested action to be performed, and the user environment are evaluated. This evaluation results in a list of user permissions granted, the cost for the action to be performed, and any restrictions.
- Rights management functions apply to content of any form, and value-added services as well.

- IBM will work closely with industry leaders, including software development companies, content providers, industry associations and other groups to formulate the container structure and processing rules standards. The goal of the IBM infoMarket Cryptolope implementation is to be fully interoperable with other systems and clearing centers implementing these standards.
- Content is made easily available by:
 - Other IBM infoMarket service facilities assist in locating appropriate content through source selection and simultaneously searching the information repositories of multiple content providers.
 - Distribution of information (that is, Cryptolope containers) can be outside the IBM infoMarket service (on other information networks, diskettes, and so forth).
- Consumers will be able to use IBM infoMarket documents in a disconnected state, that is, while not connected to a network. Over time, the IBM infoMarket service will use secure payment cards, smart cards, and other state-of-the-art technologies to support digital cash payments and other advanced functions.

For more information about the IBM infoMarket Rights Management Architecture see Appendix C, “IBM infoMarket Rights Management Architecture” on page 577.

12.1.5 Superdistribution

Since the advent of the photocopier, publishers have been dealing with people who distribute their information without consideration for the publisher′s copyright. Digital information is particularly difficult to protect. However, if you could keep track of who is using your information and how they are using it, this becomes a cost-effective and efficient way to distribute and redistribute your information.

That is Superdistribution, every customer becomes a marketer for your business.

Let′s say you have just published a report showing the results of a 20-year study of the price cycles of high tech stocks. It′s a great report and everyone should want one. Now you are faced with either selling it in a store or by mail order for, say $50.00 per copy. If you aren′t a large publisher, you may not be able to get your report on the shelves. Trying to mail a brochure to every potential customer in the world would also be quite a challenge.

It all seems pretty bleak, doesn′t it?

If you put your report up on the Web, you′ve already gotten around those two limitations. By being involved with IBM infoMarket, you′ll actually get paid for it. You could even distribute your report directly in a Cryptolope envelope using e-mail and be sure anyone who reads the document pays for it too.

Now this is where it all starts to add up.

1. Six hundred people download the report.

2. Each of them pays $10.00 to read it. IBM infoMarket will charge a percentage to handle the transaction, but there are no fees if no one buys. We handle the encryption, delivery, billing and collection for you.
3. You've just made a lot of money.

Some 450 of those people thought that the report was dynamite.

1. Each of them pass the Cryptolope envelope on to three of their colleagues.
2. Each of those 1,350 people pay another $10.00, or less, if you choose, to read it. IBM infoMarket will charge a percentage to handle the transaction.
3. You've got another pile of money in your coffers.

If only a fraction of those people pass it on, your report is being redistributed and the money keeps rolling in without any additional outlay on your part.

To your customers, this is a terrific bargain. Now they don't have to track down your report in a bookstore or library or wait for it to come in the mail. They have it in their hands immediately for a fraction of what they would have paid through more traditional channels

. For more information about Superdistribution, please see `http://www.infomarket.ibm.com`.

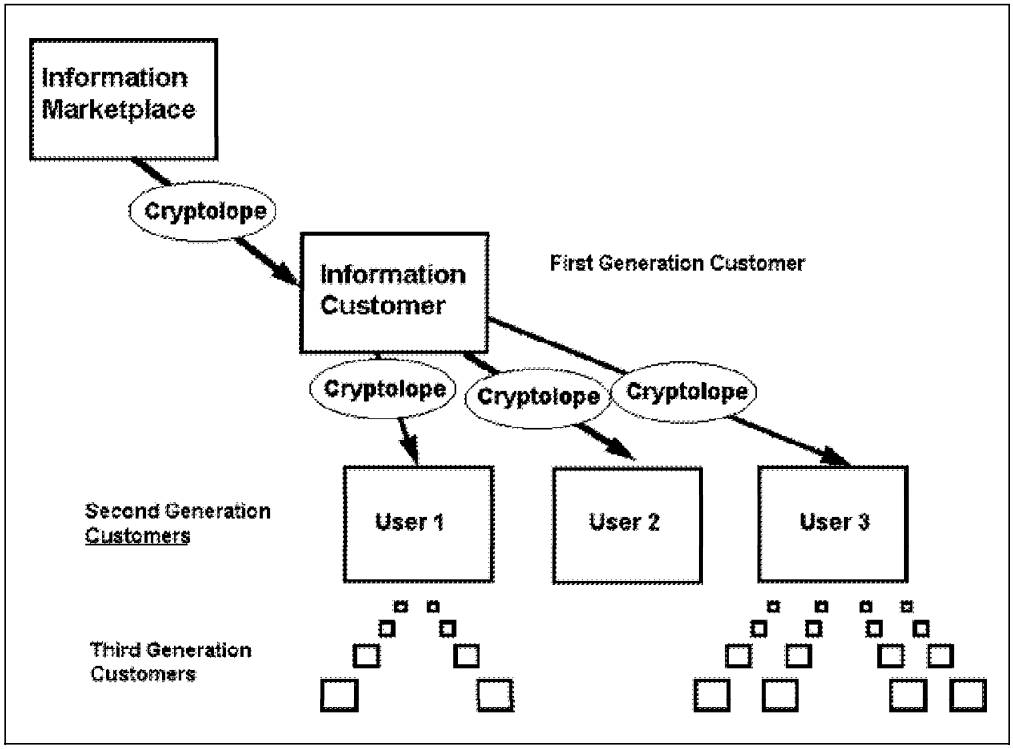


Figure 247. Superdistribution

12.1.6 IBM infoMarket Applications

These are the IBM infoMarket applications.

12.1.6.1 Cryptolope Helper Application

IBM′s Cryptolope Helper Application is a unique, revolutionary way to preview and purchase articles, literature, software, and other rights-protected digital data.

The IBM Cryptolope Helper Application is necessary for you to preview and purchase text, audio, and visual information and other rights protected digital data. Working with the browser, the Cryptolope Helper Application provides a way for you to read the terms and conditions of the copyrighted material and unseal the Cryptolope document. In the Cryptolope container you could find a copy of the latest market research report, a stock tip on a new company, demographics for market segmentation and sales forecasts, legal and travel information, pictures, music, images and more.

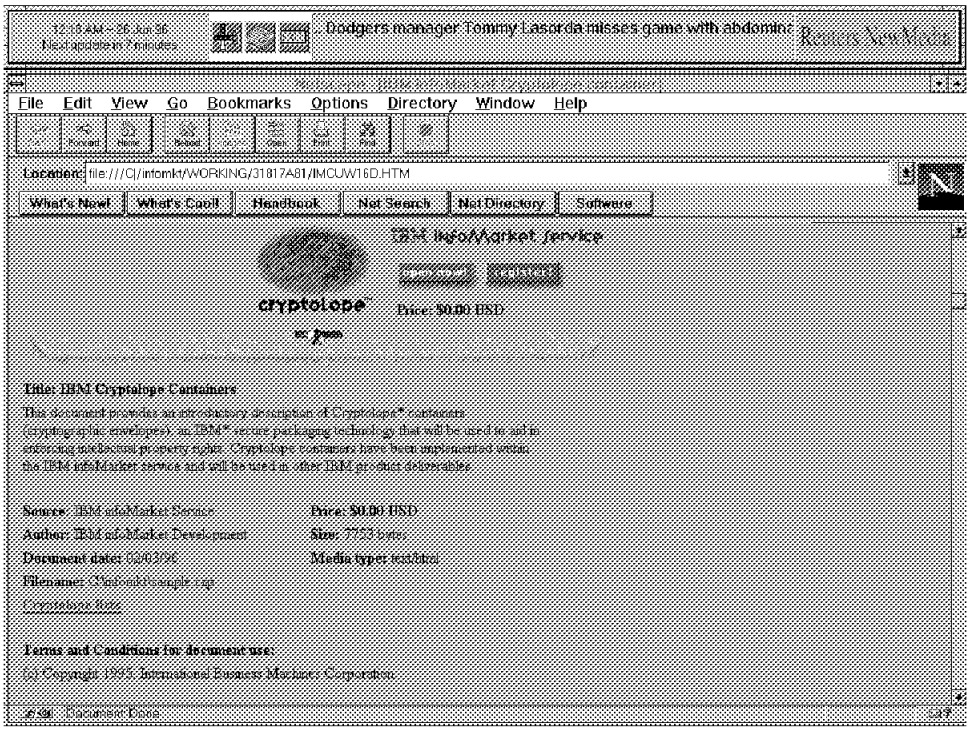


Figure 248. Cryptolope Container Viewed with the Helper

Technical Requirements:

- Windows 3.1, Windows 3.11, Windows 95, or Windows NT
- 1 MB of space on your hard drive for the Cryptolope Helper Application
- 4 MB or more of space on your hard drive for the Cryptolope containers
- 4 MB of RAM
- Netscape Navigator Version 1.2 or higher for Windows 3.x or Windows 95 or NT.

Please note: These configurations are the only supported configurations for the Cryptolope Helper Application at this time. The Cryptolope Helper

Application may run under other configurations, but we do not currently support them.

To download the Cryptolope Helper Application, route your browser to: ftp://ftp.infomkt.ibm.com/pub/cuw16v1.exe or ftp://ftp.infomkt.ibm.com/pub/cuw32v1.exe.

12.1.6.2 IBM infoMarket News Ticker

The IBM infoMarket NewsTicker represents a family of news ticker information services which is provided free to the desktops of Internet users, providing them with instant access to the most current news. Currently the NewsTicker includes news feeds from Reuters and ESPN (Sportsticker). Any Web browser can download the software required from the IBM infoMarket home page, located at www.infomarket.ibm.com. To access the service, any Windows-based Web browser can be used. Macintosh versions of the software are also planned.

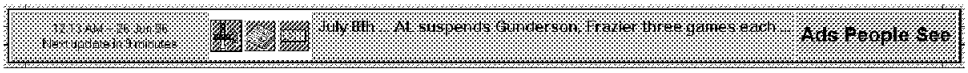


Figure 249. IBM infoMarket News Ticker

When news that you are interested in appears in the window, double-click on it. You can select the news you want to see just by clicking on the **+** or **-** icons, or you can generate the whole list.

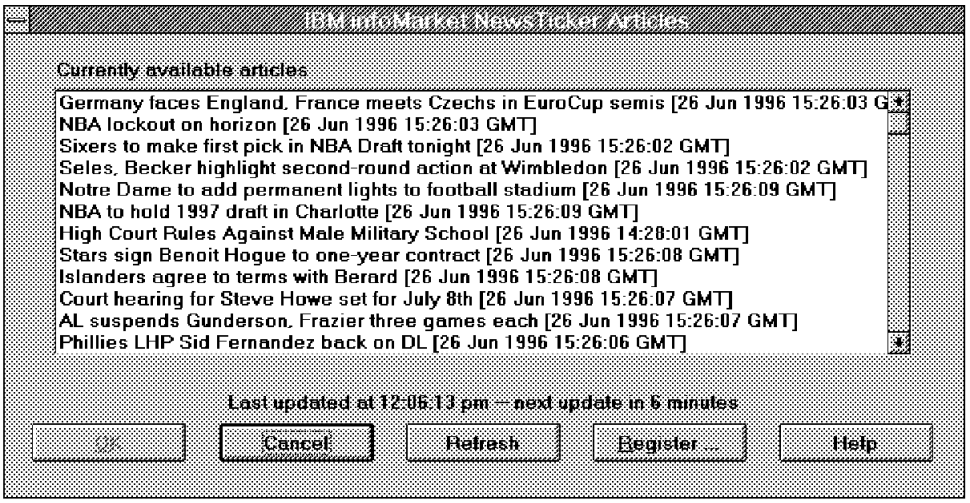


Figure 250. IBM infoMarket News Ticker Static Window

In the configuration window of the News Ticker Application you can select the news services you want to see, the time intervall to update the news and your network specifications.

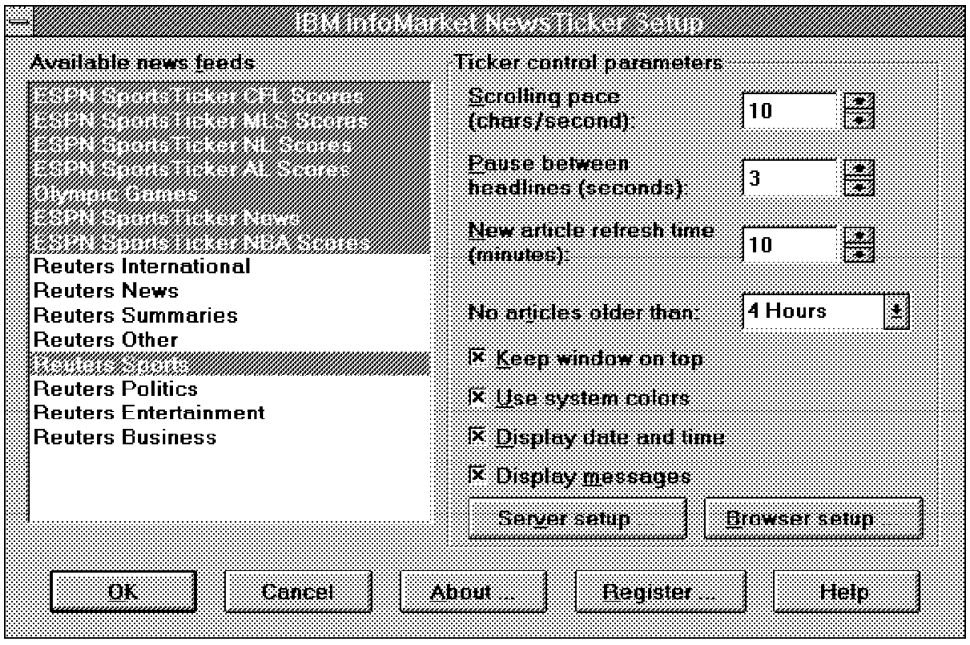


Figure 251. IBM infoMarket News Ticker Configuration Window

12.1.6.3 Plug-N-Publish Toolkits

There are several toolkits available to customize the information you want to offer.

Interface Toolkit: The Interface Toolkit does just what its name says; it is a function to develop interfaces between the customer information database and IBM infoMarket.

The Interface Toolkit also allows you to say how to offer and at what price the information should be soled.

Service Toolkit: The Service Toolkit supports you in developing further functionality such as converting of document formats or translation. The Service Toolkit also has an interface to connect the Web pages of the information owner with the IBM infoMarket Service.

Client Toolkit: Using the Client Toolkit, you are able to customize the browser to your needs. For example you can change the look to the corporate image.

For more information about the IBM infoMarket Service, see `http://www.infomarket.ibm.com`.

12.2 IBM infoSage

IBM infoSage is an information delivery system individually tailored by you to meet your specific needs. When you join, you create a detailed profile of your areas of interest. IBM infoSage uses this information as it scans a vast array of premium content resources every day to find the news items that are specifically relevant to you. This crucial information is delivered to you in a personalized newsletter twice a day either over the Web or via e-mail. And your last nine deliveries are always saved on your Web site. You can even

set the system to alert you immediately to information on topics that you deem especially vital. The IBM infoSage has the following functions included:

- Profile Editor
- Links to Related Information
- Special Editions
- Archive Search
- Stock Tracker
- Top Stories

12.2.1 Profile Editor

The IBM Profile Editor software makes it easy to create your detailed, individual interest profile. This profile is used to determine which stories should be included in your daily deliveries. Once you've registered to become a member, you can either order the Profile Editor software or download it from the Web site. With this software, you can choose and create your profile topics and adapt them to your individual needs. You do all of this offline, then the Profile Editor helps you send your completed profile to the IBM infoSage System. You only have to create your profile once, but you're free to change it as often as you like.

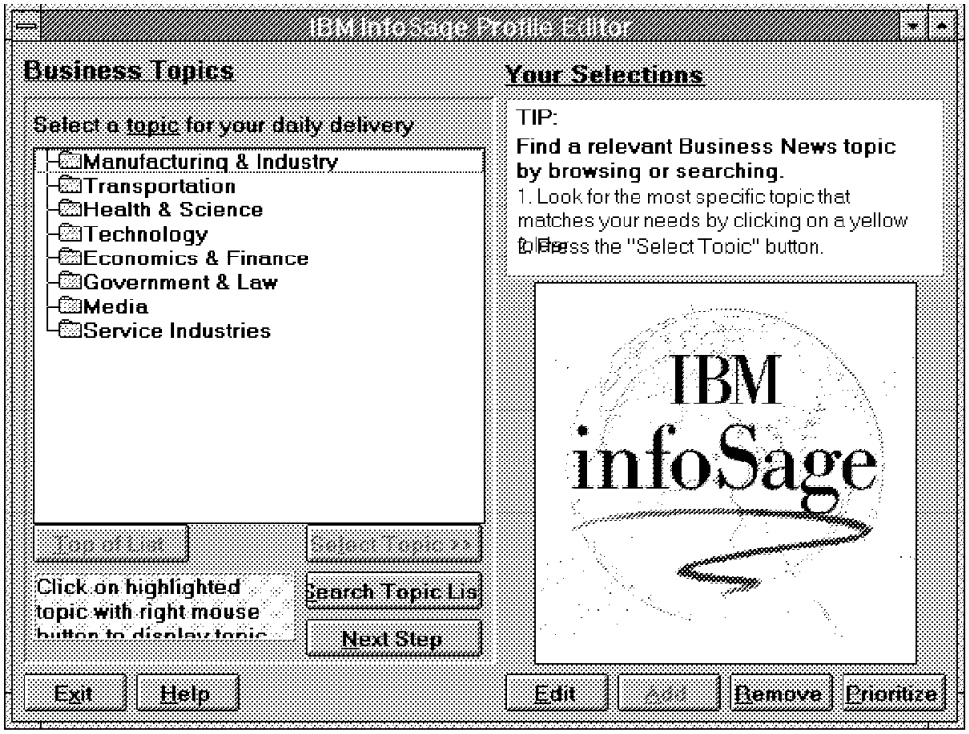


Figure 252. IBM Profile Editor

You can choose topics out of the business section and from the leisure section. After your selection of a specific topic you can personalize this topic with different keywords.

In Appendix D, “More Information about IBM infoSage” on page 581 you can see a list of the contents provided from infoSage, a few screenshots from the IBM Profile Editor and examples of results delivered through infoSage.

12.2.2 Special Editions

As an added benefit to members, IBM infoSage also offers a number of Special Edition newsletters dealing with a variety of subjects. Choose from columns on Personal Investing, Entertainment, Asset Management and many more. For your convenience, Special Editions are delivered to you via e-mail.

12.2.3 Archive Search

Another key benefit is that from your personal Web page, you have the unique ability to search the database for information you need on any topic, at any time. The IBM infoSage search works much like an Internet search engine. But, instead of searching free Web sites, our search database includes a wide array of premium content resources that are not readily available to the general public. The IBM infoSage archive search offers you a depth and breadth of information that would be difficult to find anywhere else. So you always have access to the information you need when you need it.

12.2.4 Stock Tracker

Another important function that IBM infoSage handles for you is tracking your personal stocks. You never have to squint at columns of fine print in a newspaper again to see how your investments are doing. You can specify up to 20 publicly held companies to be tracked in your personal profile. To make this information easy to access, it’s included in your morning delivery. You can also request a current stock quote (20 minute delay minimum) at any time right from your Web page.

12.2.5 Top Stories

Besides knowing what’s important to you, you probably also want to keep up on what’s important to the rest of the world. In addition to your profiled news delivery, IBM infoSage gives you a concise executive summary of each day’s top news, business and sports stories. So you can catch up on what’s going on at a glance. With IBM infoSage, the information you need comes to you. You never have to go looking for it.

12.2.6 Links

The stories in your daily deliveries are just the beginning of what the service offers. IBM infoSage also links you to even more in-depth and valuable information. These links are designed to complement the stories and articles captured by your individual profile. For instance, you might be offered detailed financial reports on a company that appears in one of your stories. These linked documents are available on a pay-per-view basis. Their prices are clearly marked before they are fulfilled. For more information or a guided tour see http://www.infosage.ibm.com.

12.3 Electronic Purchasing Service

Electronic Purchasing Service, in pilot testing at the time of writing this book links buyers electronically to their suppliers for a more efficient procurement process. The service is an advanced network-based sales and procurement solution that allows end users to locate, compare, and purchase items directly through customized electronic catalogs while providing corporate purchasing departments with better control over the process.

Electronic Purchasing Service delivers benefits to both buyers and suppliers including reduced procurement costs and greater leverage of purchasing power for buyers, and reduced marketing and order fulfillment costs for suppliers.

The IBM Electronic Purchasing Service is simple, secure and can support suppliers and buyers at all levels of technological sophistication. Suppliers work with IBM to produce custom electronic catalogs based on the agreements they have negotiated with buyers. Procurement management has browse access to the full supplier catalog. End users have access to a subset of the full catalog that contains pre-approved items and contract pricing only.

Access to supplier catalogs is password protected through Lotus Notes. Only the designated administrator can add a new user to the system. The IBM Operations Center manages the maintenance of catalog content, and facilitates the transactions between buyers and suppliers.

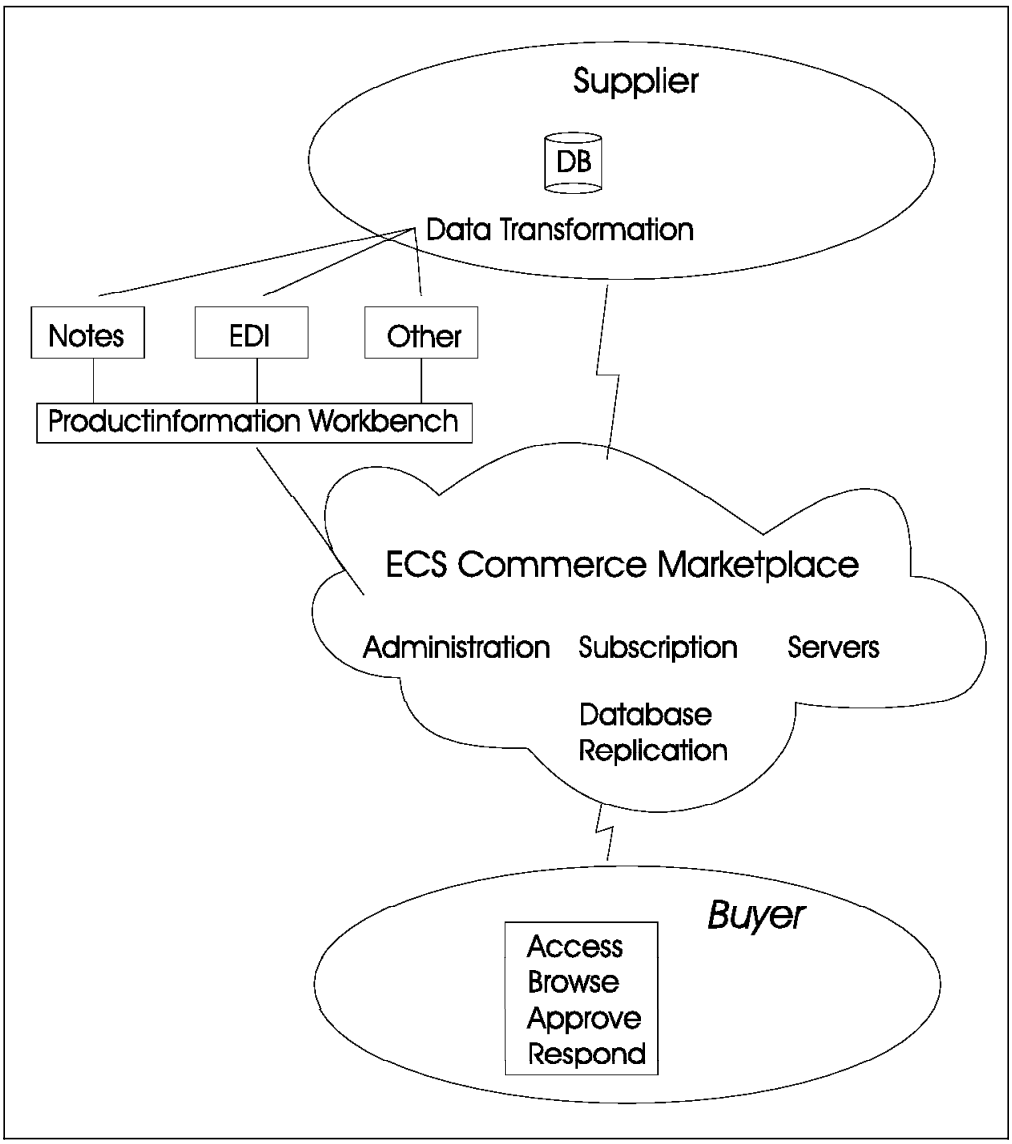


Figure 253. Electronic Commerce Exchange

Suppliers can customize catalogs to display any product with the option of including images and executables. This document attachment feature allows the supplier to provide easily viewed, detailed product information, for example specification sheets, product safety sheets, etc., that help end user make informed buying decisions.

Electronic Purchasing Service features:

- Exploits the groupware advantage of Lotus Notes
- Includes custom catalogs containing only pre-approved items/prices
- Provides catalog search capabilities
- Supports attached documents including images, text, files, or executables to augment product descriptions
- Enables timely online product and pricing updates from suppliers
- Provides the ability to save orders for future use
- Complements a re-engineered purchasing process

- Enables online browsing, selection, approvals and ordering
- Empowers end users while improving controls
- Supports three levels of electronic approvals
- Provides access to multiple supplier catalogs through a single interface
- Allows placement of custom orders, for example business cards
- Can link order payment to a purchasing card for low-cost, streamlined payment processing
- Can route POs online to three levels of approvals

For more information about the Electronic Purchasing Service, see http://www.ecs.hosting.ibm.com.

12.4 Interactive Marketing Service

The Interactive Marketing Service, in pilot at the time of writing this book enables companies to outsource their online catalogs to IBM.

The Interactive Marketing Service provides merchants with the ability to control catalog content, design and layout. Merchants can establish a personal relationship with their customers, 24 hours a day, seven days a week, using powerful search and navigation tools, intelligent agents and data mining technology. This service will be featured in a series of upcoming IBM Internet offerings for business in health care, media and other industries.

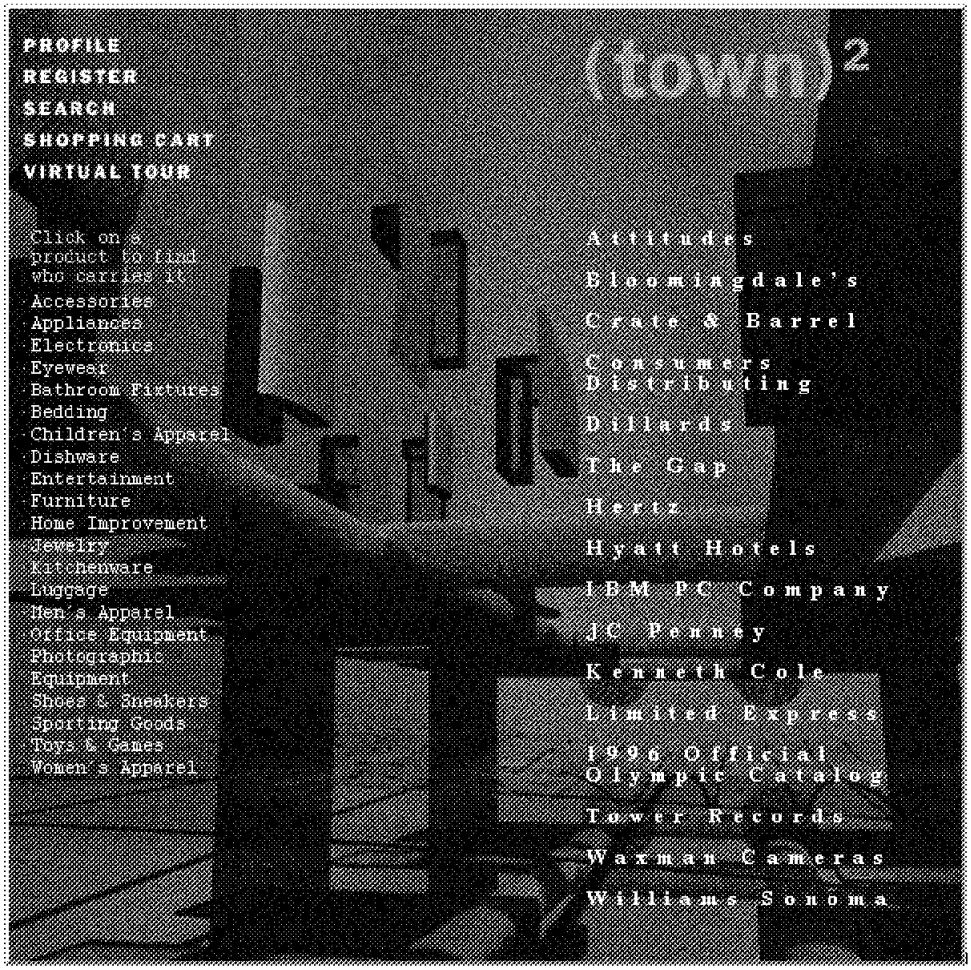


Figure 254. Sneak Preview of IBM's Shopping Site

For more information about the Interactive Marketing Service see http://www.ecs.hosting.ibm.com/ims/ims.htm/.

12.5 Net.Commerce

The Net.Commerce product allows you, as the merchant or service provider, to create an electronic store where your products or services can be sold to potential customers on the Internet's World Wide Web (WWW). Using Net.Commerce, your shoppers can browse and purchase goods and services described in your electronic store. This store will make the shoppers feel like they are shopping in a real store.

Net.Commerce can be used with a standard Web browser, such as the Netscape Navigator 2.0 or another Java-compatible browser. In addition, Lotus payment switch technology provides the integrity and the authentication necessary to allow your shoppers to securely purchase products and services over the Internet.

Net.Commerce consists of a Store Manager, a Net.Commerce director, and a Net.Commerce daemon. Figure 255 on page 538 shows these components and how they interact with other products that are part of IBM′s world of electronic commerce.

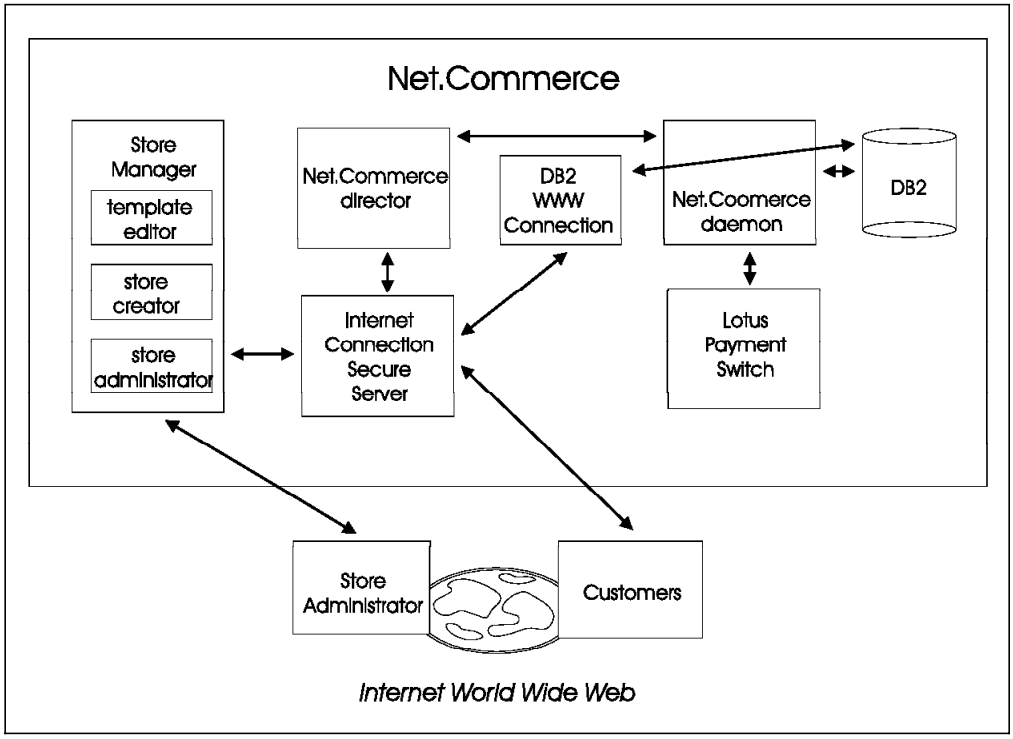


Figure 255. Net.Commerce

12.5.1 The Store Manager

Store Manager is a component of Net.Commerce that provides the tools that a store administrator needs to create and administer electronic stores. The Store Manager also provides the tools for keeping track of prices, orders, shoppers, and groups of shoppers for group discounting or group pricing.

Store Manager contains a collection of Java applets that are installed on the Net.Commerce server and which may be accessed from any Java-compatible browser on the World Wide Web. Store manager consists of the following components: the store creator, store administrator, and the template editor.

For more information about the Store Manager and its components, and how to create and maintain a virtual storefront on the World Wide Web, refer to the *Net.Commerce Store Manager Handbook.*

12.5.2 The Store Creator

The store creator is a series of easy-to-use interfaces on the World Wide Web that guide a user through the initial steps to creating a basis for an electronic store. The store creator provides the basic elements of an electronic store, and directs the user to the store administrator and to the template editor to provide the remaining content and design of the electronic store.

The store creator enables a store administrator to perform the following basic store operations:

- Create a store basis
- Configure the electronic store

- Design the store's home page
- Categorize the store's products
- Design a default store header and footer
- Design the shopping basket
- Define shopper groups
- Configure Net.Commerce

12.5.3 The Store Administrator

The store administrator is a collection of Java forms on the World Wide Web that provides easy access to entering, editing, and maintaining store information in the Merchant Server database.

Using the store administrator, a user can:

- Create an electronic store
- Configure Net.Commerce and the electronic store
- Change and maintain the stores information
- Enter and modify product and price information
- Maintain shopper records
- Maintain groups of shoppers
- Assign custom headers and footers to store pages
- Customize the store display for different shopper groups
- Keep track of orders

12.5.4 The Template Editor

The template editor provides a what-you-see-is-what-you-get (WYSIWYG) environment allowing you to design the look and feel of your electronic store, so that your shoppers feel like they are in a real store. With it you can create your store pages that includes the stores home page, interactive navigational pages and dynamic catalog pages.

12.5.5 The Net.Commerce Director

The Net.Commerce director is a non-parse header common gateway interface (pph-cgi) program allowing two-way communication between the IBM Internet Connection Secure Server and the Net.Commerce daemon. It is called by the IBM Internet Connection Secure Server to display products and services offered for sale to your shoppers. The Net.Commerce director communicates via a TCP/IP socket with the Net.Commerce daemon to quickly access the store's database. The TCP/IP communication is secured through a public/private key encryption mechanism.

12.5.6 The Net.Commerce Daemon

The Net.Commerce daemon is a program used to access information stored in a DB2 database from which your online product catalogs are built. It can assist in building pages dynamically and rapidly, in maintaining and multiplexing the connections to the database, and managing the security and administration of the Net.Commerce.

12.5.7 The Lotus Payment Switch

The Lotus payment switch performs authorization for credit card transactions when shoppers place their orders.

The transaction information is transmitted in a secure fashion to the payment server for processing. The response is returned to the Net.Commerce server where an appropriate URL tells the shopper whether the transaction has been accepted or rejected.

12.5.8 The Olympic Ticket Sale - an Example of Net.Commerce

The Atlanta 1996 Olympic Ticket Sale on the Internet is the largest electronic commerce application on the Internet at the moment. It is realized with IBM Net.Commerce. We show you with this example the possibilities of Net.Commerce.

Let's try to get some tickets!

We start at the ticket sale home page at http://sales2.atlanta.olympic.org. In the upper part of the screen you can see the heading definition done with Net.Commerce. This heading you can find on every page in the ticket sale.

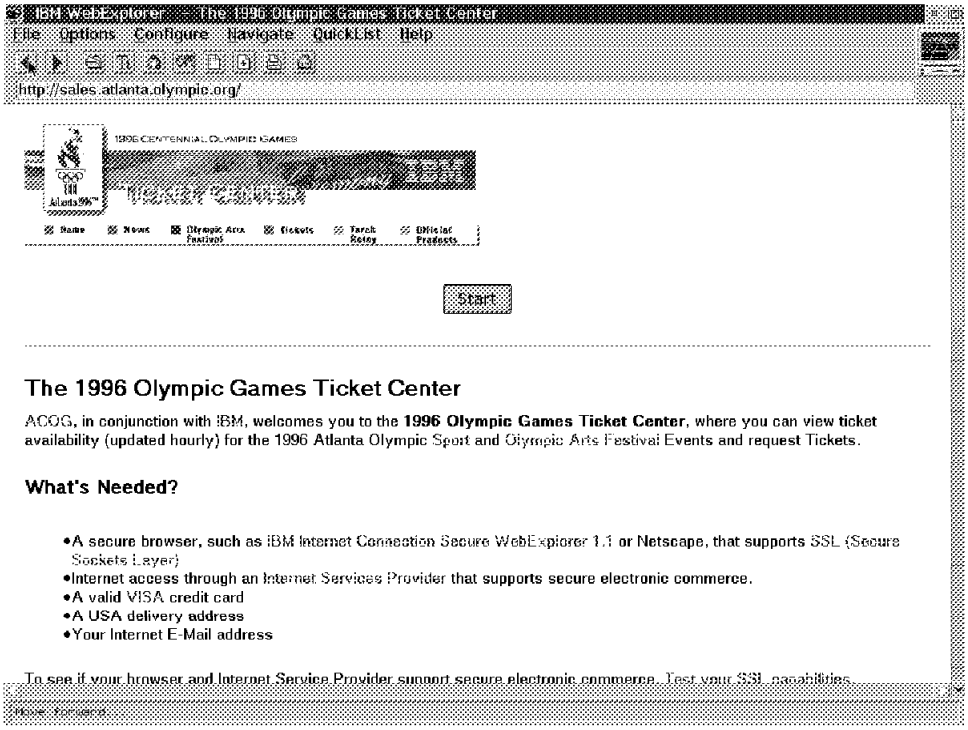


Figure 256. The Olympic Ticket Sale Start Page

After pushing the **Start** button, the selection page appears. Here you see the different search possibilities you have for getting tickets. In the same way you can build selection categories for your business using Net.Commerce.

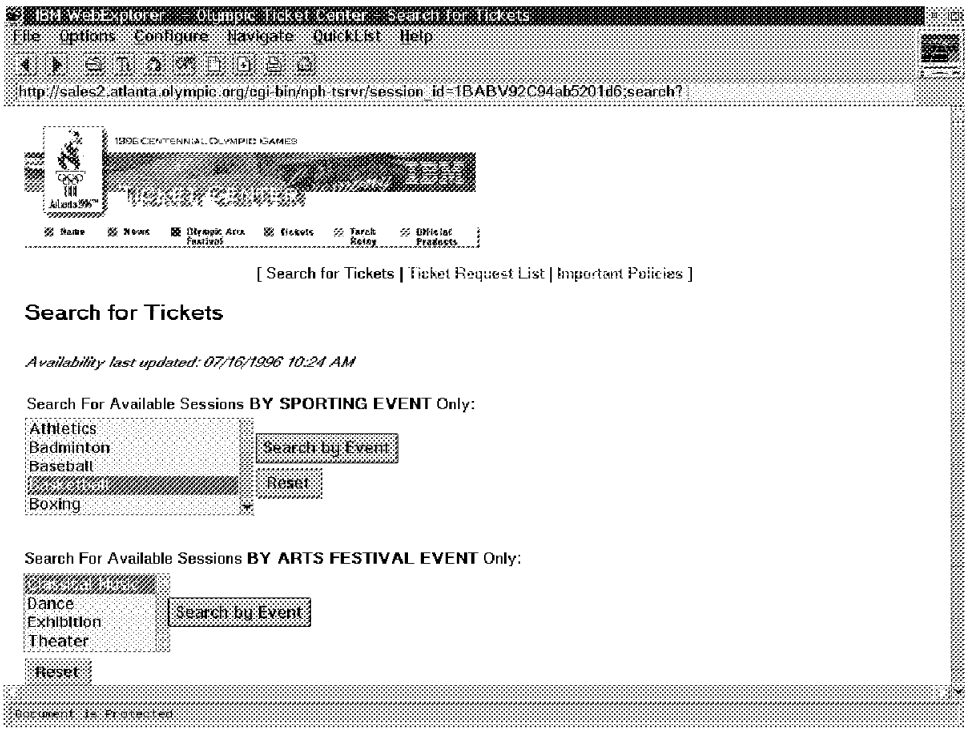


Figure 257. Search for Tickets Part 1

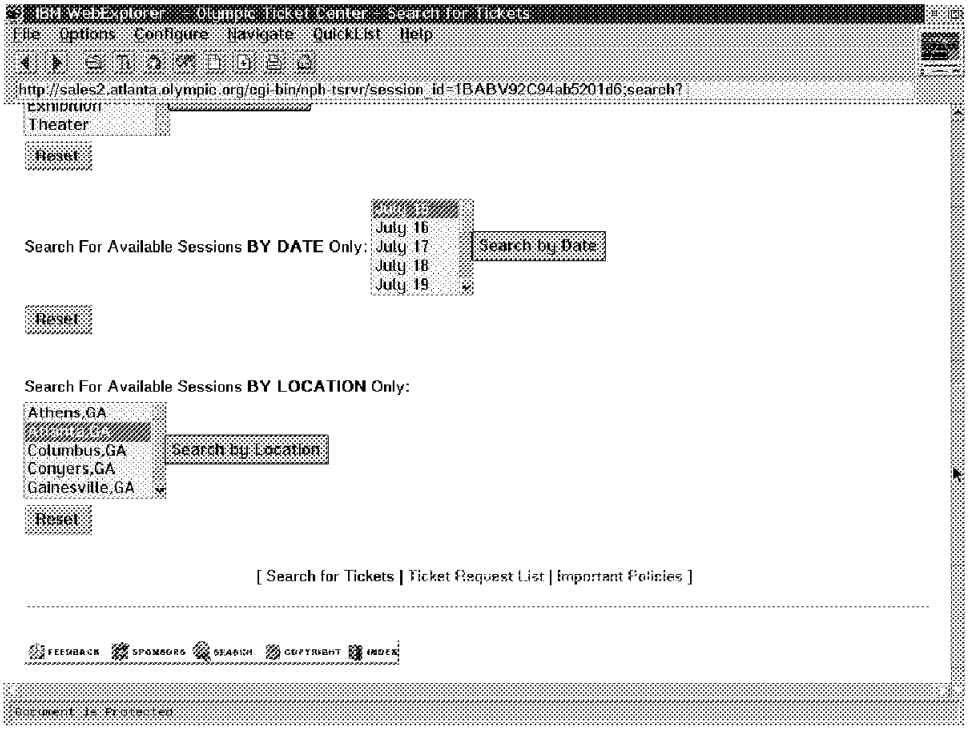


Figure 258. Search for Tickets Part 2

We want to know if there are any tickets available on the 31st of July, so we choose the **Search by Event function**. The search result showed us all events for that date where tickets were available.

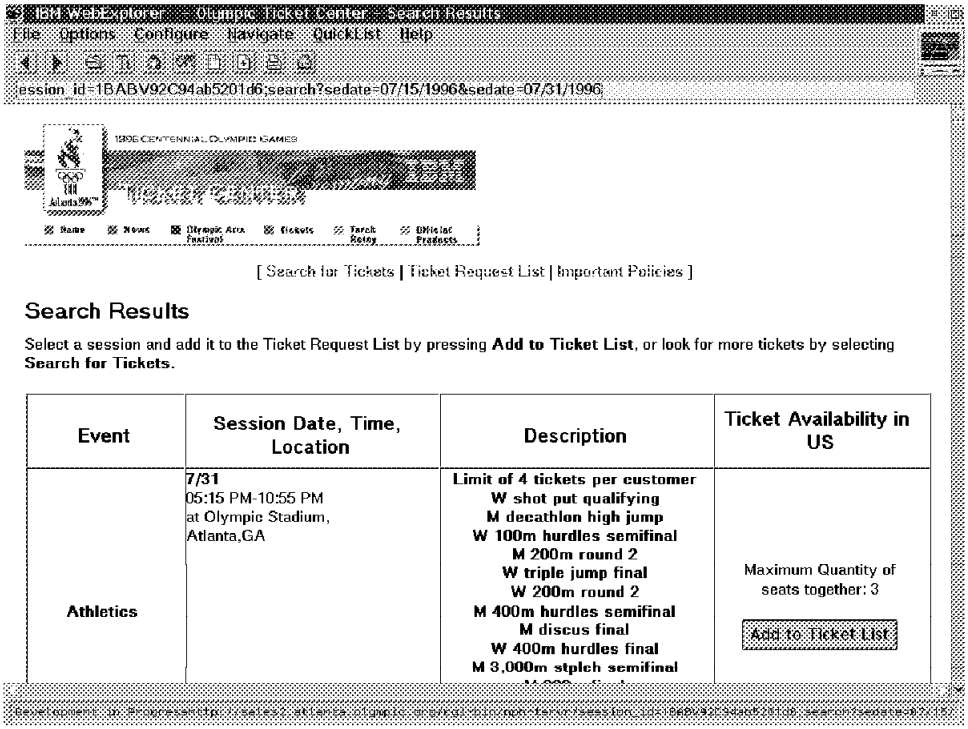


Figure 259. Result of Search by Date

We decided to go to a hockey game in the morning and to a handball game in the afternoon.

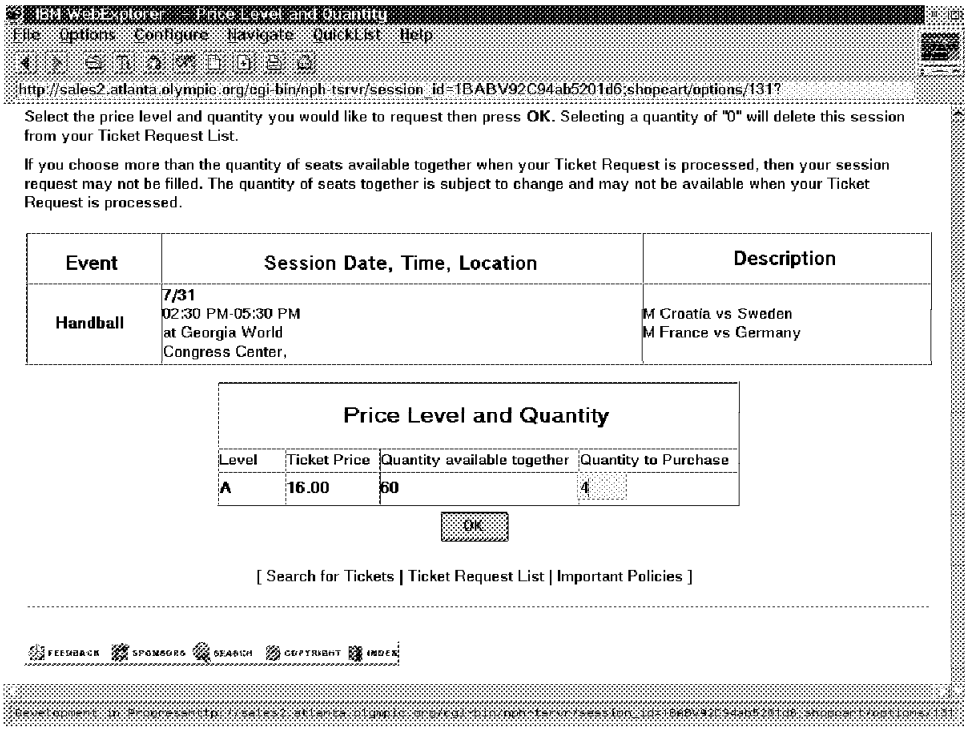


Figure 260. Ticket Price and Quantity

After every selection, we saw the list of all of our ticket requests, with the possibility to change the requests.

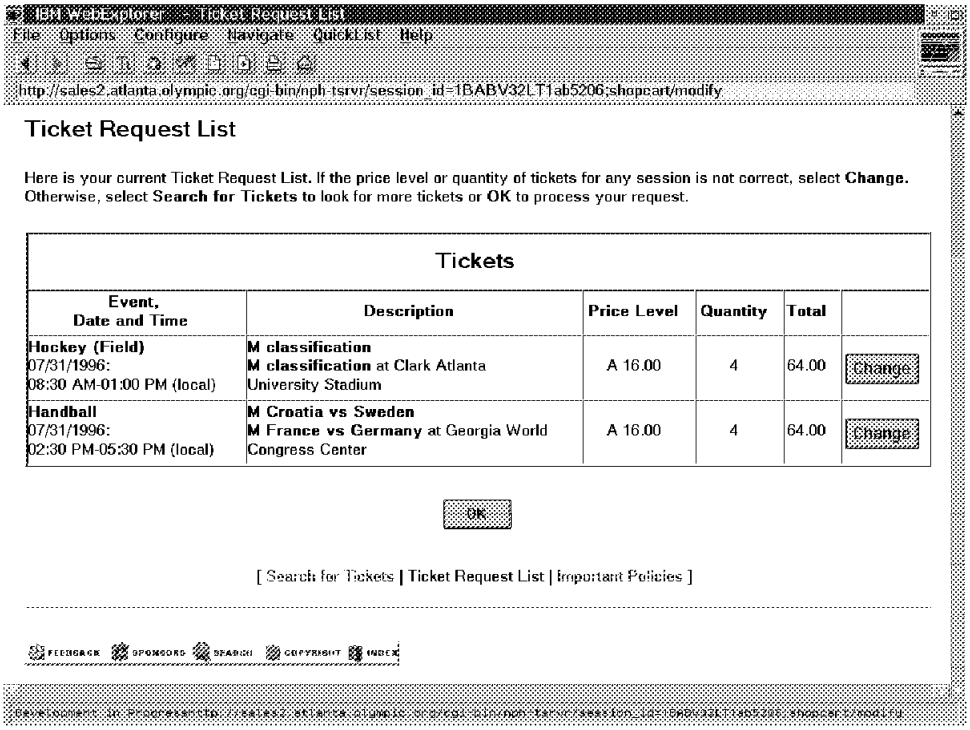


Figure 261. Ticket Request List

By selecting pushing the **OK** button in the ticket request list, we started the payment process. Net.Commerce first proves if the browser supports SSL. Our browser didn't support SSL, so we got the following page as a result:

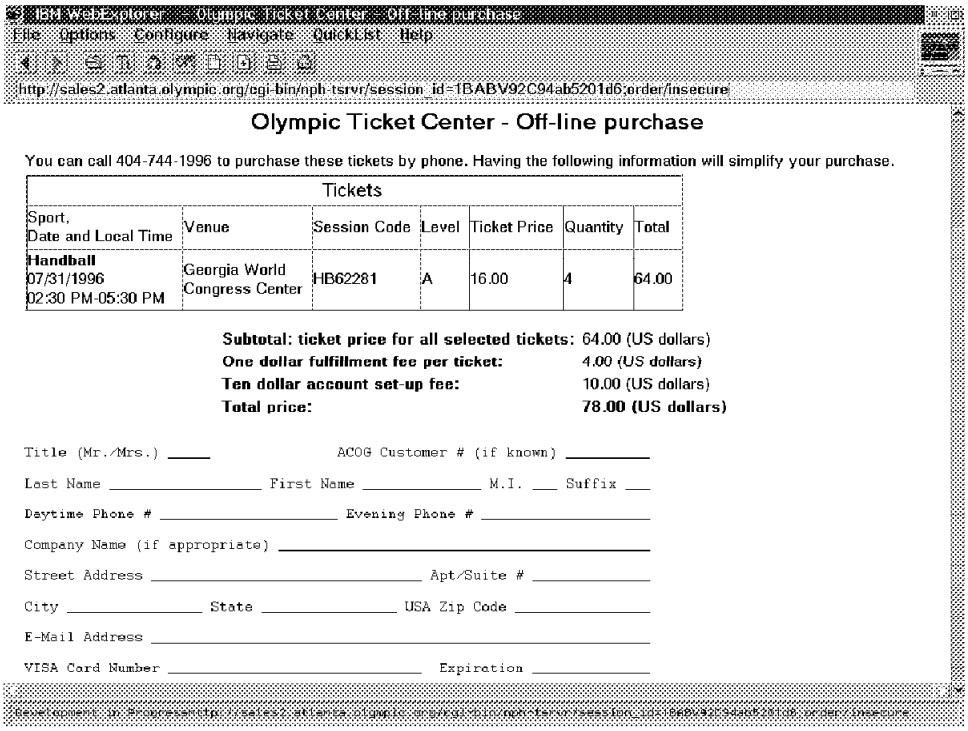


Figure 262. Unsuccessful Security Test

As you see, Net.Commerce offers your customers two ways to order and pay:

- With SSL support in your browser, your customers can order online and pay with their credit card (only VISA at the time of writing this book).
- Without SSL support they can use the Net.Commerce for selecting the products or services they want and then they can order offline.

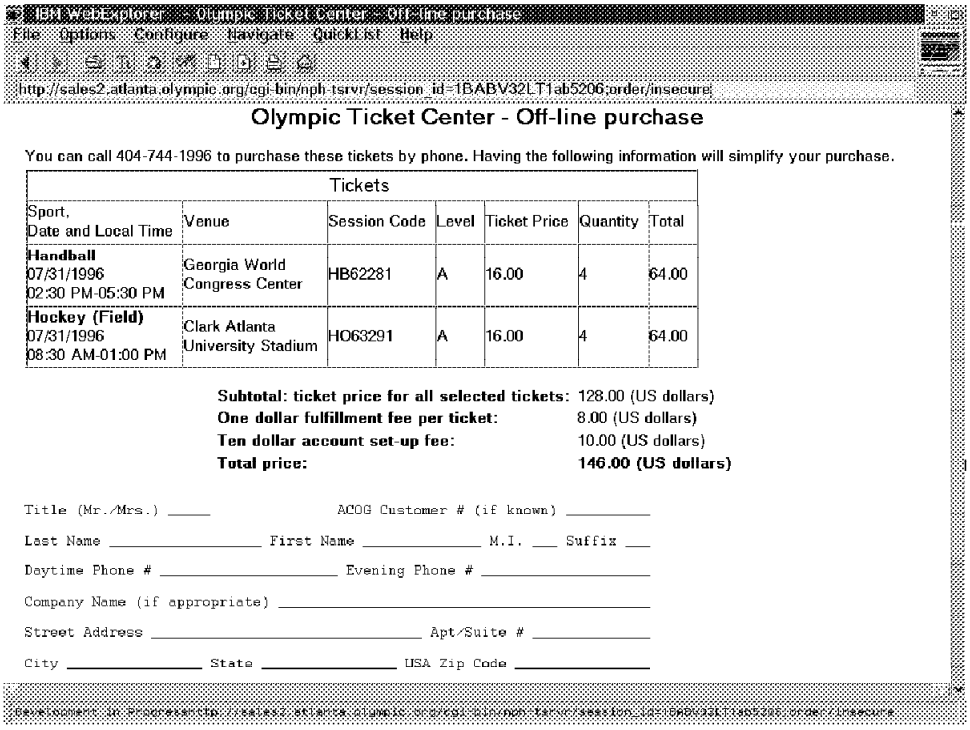


Figure 263. Offline Purchase

Chapter 13. Internet Sample Solutions

This chapter contains some sample Internet solutions.

13.1 Basic E-mail Solution

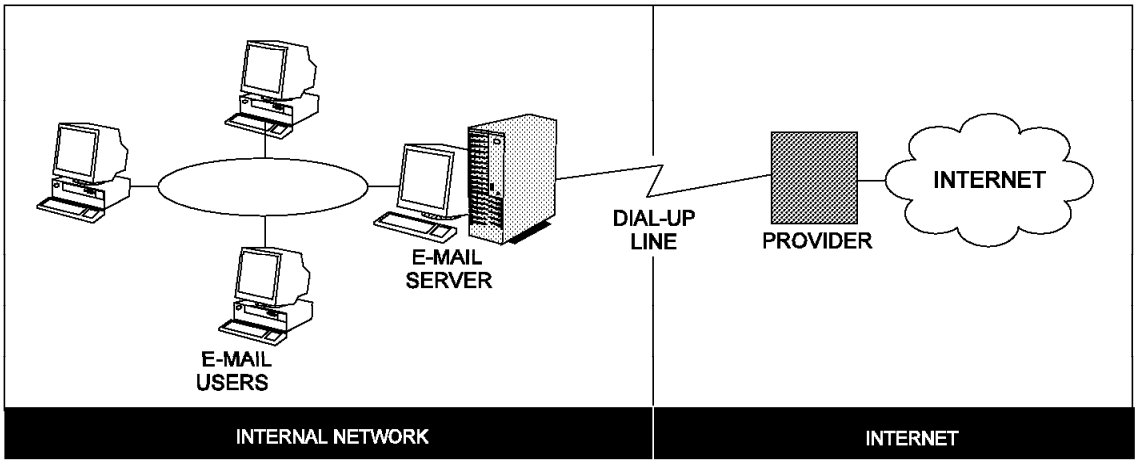


Figure 264. Basic E-mail Solution fo Small Offices

This solution uses a common phone line to connect the provider on a defined interval. The server works like a spooler holding all forward E-mail messages from the LAN, and when the provider is connected all these messages are sent to Internet, and the messages from the Internet are downloaded on the server. This kind of service must be negotiated with the service provider. Note that this solution doesn′t have an advanced security system. The configurations listed below were created considering a 25-users LAN environment.

Table 43. Basic E-mail Solution Using OS/2 Warp

Resource	Software requirements	Hardware requirements
E-mail Server	• OS/2 Warp 3.0 • SMTP and/or POP server • TCP/IP configured and running • SLIP or PPP interface configured and running	• IBM PC Server 310 • Pentium 90Mhz CPU • 32MB RAM • 2.0 GB hard disk • LAN adapter • Modem 28.8 Kbps at minimum • DAT backup tape • CR-ROM unit

© Copyright IBM Corp. 1996

Table 44. Basic E-mail solution using WindowsNT 3.5.1

Resource	Software requirements	Hardware requirements
E-mail server	• WindowsNT 3.5.1 • SMTP and/or POP server • TCP/IP configured and running • SLIP or PPP interface configured and running	• IBM PC Server 310 • Pentium 90 Mhz CPU • 32 MB RAM • 2.0 GB hard disk • LAN adapter • Modem 28.8 Kbps at minimum • DAT backup tape • CR-ROM unit

Table 45. Basic E-mail Solution Using AIX 4.1.4

Resource	Software requirements	Hardware requirements
E-mail Server	• AIX 4.1.4 • SMTP and/or POP server • TCP/IP configured and running • SLIP or PPP interface configured and running	• IBM RS/6000 Model 43P • PowerPC 100 Mhz CPU • 32 MB RAM • 3.0 GB hard disk • LAN adapter • Modem 28.8 Kbps at minimum • DAT backup tape • CR-ROM unit

Table 46. Basic E-mail Solution Using Lotus Notes

Resource	Software requirements	Hardware requirements
Workgroup, workflow and e-mail server	• OS/2 Warp 3.0 • Lotus Notes 4.0 or 4.1 for OS/2 • Lotus Notes SMTP/MIME MTA • TCP/IP configured and running • SLIP or PPP interface configured and running	• IBM PC Server 310 • Pentium 90 Mhz CPU • 32 MB RAM • 2.0 GB hard disk • LAN adapter • Modem 28.8 Kbps at minimum • DAT backup tape • CR-ROM unit

Table 47. Clients Specifications		
Resource	**Software requirements**	**Hardware requirements**
LAN e-mail client	• IBM DOS, OS/2, AIX, MS-DOS, Windows 3.x, 95 or NT • TCP/IP configured and running • LAN interface configured and running • E-mail reader/sender • Lotus Notes client (if using Lotus Notes as a server)	• IBM PC or compatible • 486DX4 or Pentium CPU • 8 MB RAM • 500 MB hard disk • LAN adapter

13.2 Corporative Secure LAN Solution

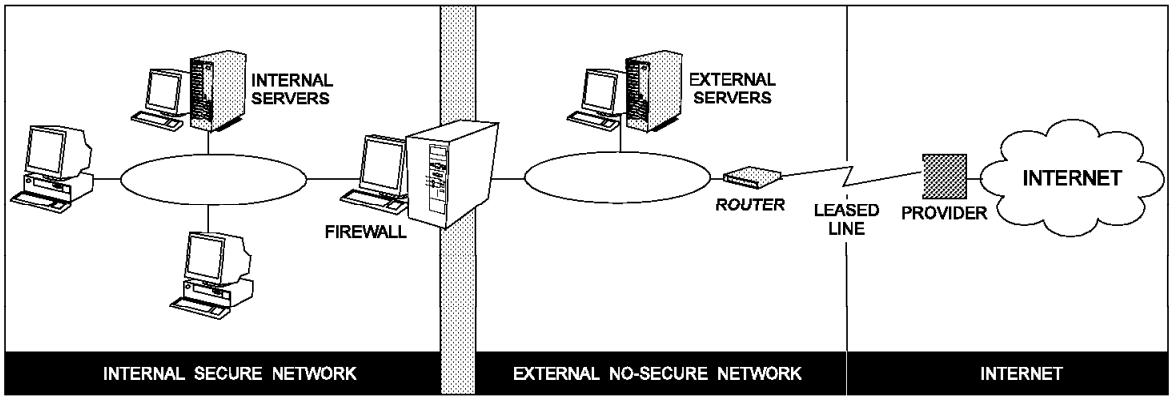


Figure 265. Corporative Secure LAN Solution. A secure way to integrate the existing LAN with the Internet.

This solution provides a relatively simple way to connect an existing LAN to the Internet. In theory, all kinds of LANs can be connected to the Internet, such as Ethernet, Fast Ethernet, ATM, token-ring and FDDI. Usually, the external LAN is a classical Ethernet 10BaseT LAN, because the bandwidth between the site and the service provider is not fast enough to justify a high-speed LAN structure.

If you have a fast Ethernet LAN installed and you are going to connect this LAN to the Internet using an IBM firewall solution you need to consider an additional bridge, because the IBM RS/6000 machines do not support fast Ethernet adapters.

All other LAN technologies are supported by the RS/6000 family, such as:

- Ethernet 10BaseT, 10Base2, 10Base5
- FDDI and CDDI
- ATM
- Token-ring

For the external servers you can use PC Servers, RS/6000, AS/400 or S/390 systems. It depends on the application and the expected performance. If you are using a low speed connection, such as 56 kbps or 64 kbps you don't need a high performance machine, because the link will be a restriction on the data flow.

Table 48. Corporative Secure Solution Specifications

Resource	Software requirements	Hardware requirements
Firewall	• AIX 4.1.4 • IBM Secure Network Gateway for AIX • Two LAN interfaces configured and running	• IBM RS/6000 Model 43P • PowerPC 133 Mhz CPU • 32 MB RAM • 2.0 GB hard disk • Two LAN adapters
External network		Ethernet 10BaseT recommended, using IBM 8222 or IBM 8224 hubs
External servers		Depending on the service that will be provided, like WWW, FTP, e-mail, CHAT, etc.
Router	IP routing support level	• IBM 2210 Model 12E • 4 MB RAM
Leased line		You can use microwave radio, satellite, common leased-lines, ISDN, etc.
Provider		IBM Global Network servives

Table 49. Clients Specifications

Resource	Software requirements	Hardware requirements
LAN client	• IBM DOS, OS/2, AIX, MS-DOS, Windows 3.x, 95 or NT • TCP/IP configured and running • LAN interface configured and running • Browser compatible with the operating system	• IBM PC or compatible • 486DX4 or Pentium CPU • 8 MB RAM • 500 MB hard disk • LAN adapter

13.3 Electronic Commerce Solution

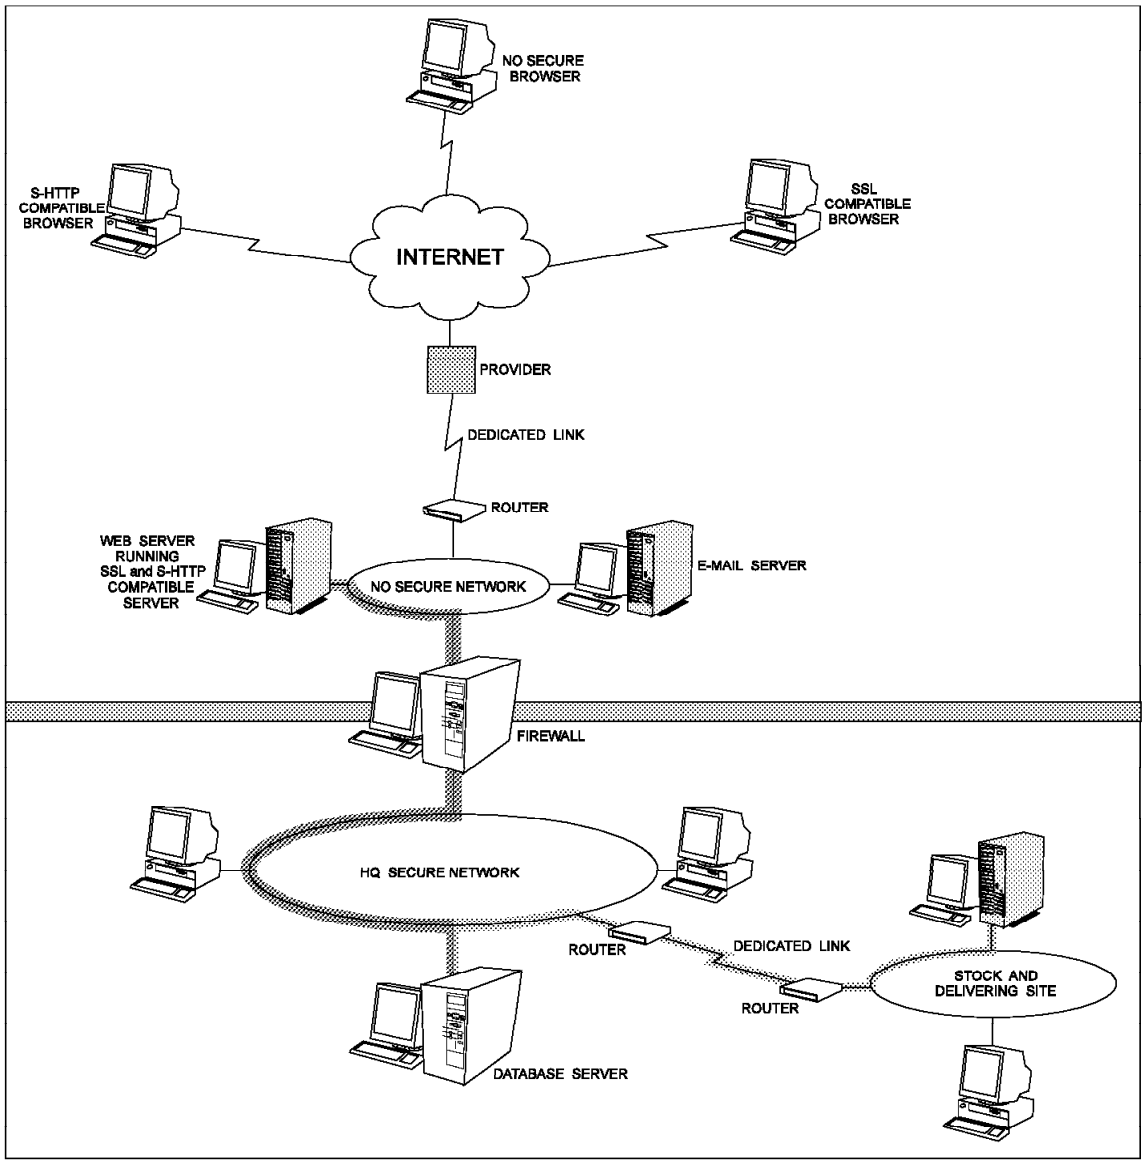


Figure 266. Electronic Commerce Solution. Electronic sales envirommment with built-in secure resources.

The solution showed in the drawing above is a basic electronic commerce solution. You can add more features on this solution providing more resources and a better service to the customers.

There are some very important considerations about this solution, such as:

- **Link bandwidth:** You need to use a link that provide a good response time to the customers.
- **Server performance:** The server performance is directly related with the link bandwidth. Always choose servers that can receive upgrades on the storage capacity, memory and if possible, on processors.
- **Security:** You must develop applications that uses all security transaction technologies available, such as S-HTTP, SSL and e-money. If you have a site that uses these technologies you are able to provide a good service for all kind of customers using all kinds of browsers.

- **Database server:** This is a vital server where all information about products availability, customers information, prices, etc. will be stored. Always look for upgradeable servers. Be careful when you choose a database software. Some databases have limitation s when used with Web integrated enviromment. The IBM Web servers can be easily integrated with DB/2 servers running on OS/2, WindowsNT, RS/6000, AS/400 and mainframes. The IBM servers also support CICS integration.
- **Firewall:** The firewall is a vital equipment on this solution, because it will provide the security for the internal LAN and to the internal servers, such as the database server.

You can connect the ″headquarter″ LAN, where are all servers, with remote LANs on stock and delivering sites, providing a quick service to the customers and a real efficient logistic enviromment.

All computers on the internal LAN will be able to access the Internet using all resources, such as e-mail, WWW, Gopher, FTP, Telnet, etc.

Table 50 (Page 1 of 2). Electronic Commerce Solution Specifications

Resource	Software requirements	Hardware requirements
Firewall	• AIX 4.1.4 • IBM Secure Network Gateway for AIX • Two LAN interfaces configured and running	• IBM RS/6000 Model 43P • PowerPC 133 Mhz CPU • 64 MB RAM • 4.0 GB hard disk • Two LAN adapters
External network		Ethernet 10BaseT recommended, using IBM 8222 or IBM 8224 hubs
Option #1 - WindowsNT server	• WindowsNT 3.5.1 or later • IBM Internet Connection Secure Server • IBM Net.Commerce Server for WindowsNT • IBM WWW DB/2 Gateway for WindowsNT • TCP/IP configured and running • LAN interface configured and running • MS-Internet Explorer or Netscape Navigator 2.0	• IBM PC Server 310 • Pentium 90Mhz CPU • 32 MB RAM • 2.0 GB hard disk • LAN adapter • DAT backup tape • CR-ROM unit

Table 50 (Page 2 of 2). Electronic Commerce Solution Specifications

Resource	Software requirements	Hardware requirements
Option #2 - IBM AIX server	• IBM AIX 4.1 or later • IBM Internet Connection Secure Server • IBM Net.Commerce Server for WindowsNT • IBM WWW DB/2 Gateway for WindowsNT • TCP/IP configured and running • LAN interface configured and running • IBM WebExplorer or Netscape Navigator 2.0	• IBM RS/6000 Model C10 • PowerPC 120 Mhz CPU • 64 MB RAM • 4.0 GB hard disk • LAN adapter • DAT backup tape • CR-ROM unit
Database server	• IBM AIX 4.1 or later • IBM DB/2 Database server for AIX • TCP/IP configured and running • LAN interface configured and running	• IBM RS/6000 Model C10 • PowerPC 120 Mhz CPU • 64 MB RAM • 6.0 GB hard disk • LAN adapter • DAT backup tape • CR-ROM unit
Router	IP routing support level	• IBM 2210 Model 12E • 8MB RAM
Leased line		You can use microwave radio, satellite, common leased-lines, ISDN, etc. The minimum recommended link speed is 128Kbp/s
Provider		IBM Global Network servives

Table 51. Clients Specifications on the Internal LAN

Resource	Software requirements	Hardware requirements
LAN client	• IBM DOS, OS/2, AIX, MS-DOS, Windows 3.x, 95 or NT • TCP/IP configured and running • LAN interface configured and running • Browser compatible with the operating system	• IBM PC or compatible • 486DX4 or Pentium CPU • 8 MB RAM • 500 MB hard disk • LAN adapter

Chapter 14. Consulting Services

IBM′s consultancy is primarily provided by IBM Consulting Group.

IBM Consulting Group provides management and information technology (I/T) consulting services to corporations and organizations worldwide. Business transformation projects performed for clients range from reengineering to redefining the business, and often examining all aspects of a firm′s operations, such as organizational structure, processes and resources.

IBM Consulting Group is committed to helping clients gain maximum value from their technology investments, focusing on how the technology, whether from IBM or other sources, can best be aligned with business strategy. The group has the capability and experience to help create an overall I/T blueprint, increasing application development and operations effectiveness, as well as architecture and design.

For additional information about IBM Consulting, refer to:

- http://www.consult.ibm.com

14.1 Management Information Technology Consulting Service Lines

Five key service lines provide clients with access to IBM Consulting Group′s core expertise in key areas of management and I/T consulting. Service lines include:

- Transformation Services
- I/T Consulting Services
- Integration Services
- General Business
- Object Technology Services

14.1.1 Transformation Services

The transformation services are:

Business Transformation

Assists clients in reengineering efforts by helping senior management create and execute a transformation plan to increase profitability, improve customer service and quality and reduce product and service development cycles and costs.

I/T Strategy

Helps organizations ensure that their I/T strategies foster business success rather than constrain it (I/T strategies are aligned with business strategies).

Management Technologies

Applies advanced analytical techniques, business modeling, simulation and optimization technologies to solve complex business problems.

© Copyright IBM Corp. 1996

Workflow Management

Works with clients to redesign the flow of documents using imaging technology to develop the most effective operational environment. Imaging technology, when integrated successfully with I/T, can help reduce costs, control the work flow and result in a competitive advantage.

14.1.2 I/T Consulting Services

Application Development Effectiveness

Helps clients assess and improve their application development processes. Areas considered are the effective use of AD methods, organizational issues, teaming and team building, skills assessment and how well projects are managed. The intent is to enable the client′s I/S team to be more effective in addressing their end user′s applications needs.

Business Recovery

Assists clients with determining and managing their business protection and recovery program. Provides a broad set of consultant skills that focus on risk management and disaster avoidance, total enterprise-wide recovery strategies, and recovery plan development and implementation. Engagements supported are cross industry and include all platforms.

Information Systems Management

Assists clients in managing enterprise-wide information systems environments. Helps them to take advantage of systems management processes and disciplines across all I/T platforms, including client/server, network and data center systems.

Networking

Assists clients in developing network strategies and creating network architectures and designs that meet strategic business goals. Skills encompass both current and emerging technologies including high-speed switched networking; Internet/intranets; client/server, distributed and megacenter; public and private wide area networks; local area networks; voice, data, multimedia, and video.

I/T Planning

Assists in planning architectures for applications, data and technology, and developing implementation plans that enhance clients′ competitive position.

14.1.3 Integration Services

Life Cycle

Provides consulting services relative to the study, design and implementation of application solutions to client business problems. These services include assistance with the following application-related activities: requirement studies, architectures, data/processing modeling, technology selection, application integration and application design, build, test and roll out.

Rapid Solutions

Works with client end users to develop quick, proof-of-concept application solution models, which can be deployed in unusually short time frames to reduce the cost of systems development. Rapid solutions are usually workstation based.

Redevelopment

Consults with clients about redevelopment of legacy applications for client/server systems.

Systems Integration

Provides a full range of services and offerings for the implementation of integrated I/T solutions that support the client's business needs.

14.1.4 General Business

Provides a broad set of management consulting skills and capabilities to intermediate-sized businesses, emphasizing a holistic approach to align and balance the enterprise's primary processes, management and control systems, and information systems.

14.1.5 Technology Services

Provides consulting in object-oriented technologies to help clients speed the application development process by adapting to the new Object Technology (OT) environment. Helps clients train mainframe programmers to use OT in the desktop environment.

14.2 Industry Specializations

The IBM Consulting Group personnel with extensive strategic, organizational and cultural knowledge tailor services to each client's unique business and competitive environment.

- **Distribution**:

 Retail and wholesale distribution companies.

- **Finance**:

 Banks and securities firms.

- **Government**:

 Federal, state and local government entities.

- **Healthcare**:

 Healthcare payers, providers, suppliers and pharmaceutical companies.

- **Higher Education**:

 Colleges and universities.

- **Insurance**:

 Property and casualty, and life and health insurance providers.

- **Manufacturing**:

Manufacturers (such as auto, aerospace, electronics) and process companies (chemical, pharmaceutical, etc.).

- **Petroleum**:

 Worldwide integrated oil companies, large natural gas transmission companies, integrated national oil companies, and large non-integrated oil and gas companies.

- **Telecommunications & Media**:

 Telephone, cable and wireless, entertainment, broadcasting, printing and media, publishing companies.

- **Transportation**:

 Companies that move freight by rail, roads or water.

- **Travel**:

 Major airlines, lodging and travel agents, airline reservation systems, casinos and car rental agencies. Utilities: Electric and gas utilities, independent power producers, water treatment and waste management companies.

- **Cross Industry**:

 Develops client solutions in areas such as work management, image, multimedia, GIS, scientific/technical solutions, as well as emerging application areas such as sales force automation.

14.3 Internet Consulting and Services

IBM provides comprehensive consulting and services to get clients′ businesses up and running on the Internet, quickly and securely. IBM is committed to a multivendor, open systems environments with services that span the full range of Internet requirements including:

- Understanding how to leverage Internet business applications for competitive advantage
- Formulating an Internet strategy that supports client′s business plan
- Enabling a secure Internet connection to protect sensitive corporate information
- Developing Web applications and integrating them into existing business systems
- Building corporate intranets
- Developing network based applications

IBM provides all the expertise you need to get the Internet working for your business.

Business Transformation

For many companies, the Internet offers an unlimited opportunity with respect to interactions with customers, suppliers and business partners. With knowledge and experience across dozens of industries, IBM′s business consultants can help clients to understand how the Internet can be used to gain competitive advantage, determine the key business process changes required, and identify the organizational impacts of these changes.

Strategy and Assessment

Our Information Technology (I/T) consultants can assist in determining the best way to leverage and deploy the Internet in support of client′s business strategy. Areas IBM covers include business case, development, technical and content readiness assessment, cost and benefit assessment, evaluation of sourcing alternatives, identification and management of risk, and impact of an Internet solution on client′s current business network. In addition, IBM can design and implement I/T infrastructures that support secure Internet and intranet solutions.

Web Application Development

Using the latest Internet and multimedia software and technologies, IBM′s experts can design and construct state-of-the-art Internet and intranet business applications. IBM′s services encompass planning and requirements definition, graphic design, graphics creation, video production and post production, multimedia integration, data conversion and migration, custom Web scripting, testing, implementation, project management and technical support.

In addition, IBM Global Network Content Services, provides an outsourcing alternative for your Web content. Clients′ Web applications can be hosted on an IBM Global Network server in a production environment that includes hardware, software and a high-bandwidth Internet link. For additional information about IBM Global Network Content Services, refer to Chapter 11, “Content Services on the Internet” on page 451.

Implementation

IBM can help install the hardware and software required to quickly enable a secure Internet connection. These services include turn key implementation of Internet services such as e-mail, FTP, Telnet, USENET, and Web services.

Through IBM′s family of services, IBM can quickly install and establish a Web server in the client′s environment. These services streamline the implementation of a Web server and include planning, software installation and configuration, creation of an initial home page, project management and basic training.

Internet Implementation Integrated Solution Offering helps customers quickly enable a corporate Internet connection. It is designed for clients who need assistance in access to get connected to the Internet , and presence to make information available through the Internet. The Integrated Solution Offering (ISO) is jointly owned and operated by ISSC Cross Industry and ISSC Consulting and Services, Boston.

Specific customer requirements addressed include:

- Determining how ready the business is for an Internet connection
- Setting up a network infrastructure to support an Internet connection
- Establishing a presence on the World Wide Web

A complete flexible range of services are available:

- Internet Readiness Assessment

- TCP/IP integration
- E-mail integration
- World Wide Web servers
- Content design, construction, testing
- Online directory assistance
- List servers, News servers, Gopher servers
- Anonymous FTP
- UNIX Shell Accounts

Security

IBM can assist clients in creating and maintaining a secure environment for their critical business information and systems. IBM provides the technology, consulting and services necessary to help clients assess, manage, contain, and prevent potential system and network security problems.

For additional information about Security, refer to:

- Chapter 8, “Security on the Internet” on page 339
- http://www.ibm.com/security

Appendix A. The IAB

This appendix contains information that can be a complement to understanding the administrative framework on which SNMP and the Internet community in general rely on. The IAB structure and RFCs are the main topics.

A.1 The Internet Activities Board (IAB)

As Internet research activity increased during the 1970s, it was necessary to establish an informal committee to provide technical guidance for the evolution of the protocol suite. In 1979 a group called the Internet Configuration Control Board (ICCB) was established.

In 1983, the Defense Communications Agency (DCA) declared the TCP/IP protocol suite to be standard for the ARPANET, and the ICCB was reorganized. The reorganized group was called the Internet Activities Board (IAB).

The IAB is the coordinating committee for Internet design, engineering and management. It is formed by researchers and professionals with an interest in the development of the Internet. The IAB focuses on the TCP/IP protocol suite and extensions to the Internet system to support multiple protocol suites. All IAB members are required to have at least one other major role in the Internet community in addition to their IAB membership. The IAB has a chairman that serves a term of two years. New members are appointed by the chairman of the IAB with the advice and consent of the remaining IAB members.

The IAB has the following two primary subsidiary task forces:

1. Internet Engineering Task Force (IETF)
2. Internet Research Task Force (IRTF)

Each of these task forces is led by a chairman and guided by a Steering Group.

The IETF focuses on short and mid-term protocol and architectural issues to make the Internet function properly. The IETF is a large open community of network designers, operators, vendors, and researchers, divided into eight technical areas, each with its own director. Each area has its own working groups to explore situations. The IETF chairman and the eight area directors make up the Internet Engineering steering group (IESG).

The IRTF focuses on research of TCP/IP protocols and architecture. It is formed by a community of network researchers. The IRTF is formed by a set of research groups (RGs), each focusing on a broad area of research. The IRTF chairman and each of the RG chairs make up the Internet Research Steering Group (IRSG).

In the area of network protocols, the distinction between research and engineering is not always clear. Thus, it is not unusual that IETF and IRTF activities overlap. Membership overlap between the two task forces is considerable and is considered vital for cross-fertilization and technology transfer.

© Copyright IBM Corp. 1996

A.1.1 Request for Comments (RFC)

A Request for Comments (RFC) is the principal vehicle by which IAB decisions are propagated to the Internet community. The RFCs are a series of notes which were initiated in 1969 as a means for documenting the development of the original ARPANET protocol suite (RFC1000). Most RFCs are intended to promote comments and discussion, although a small proportion of RFCs document Internet standards. These in particular are marked in a status section to indicate one of required, recommended, elective, limited use, or not recommended (see Table 53 on page 562). An RFC summarizing the status of all standard RFCs is published regularly (RFC1100).

Each RFC has a number assigned to it by the RFC editor who is a member of the IAB. Each time an existing RFC text is revised, a new RFC number is assigned. The new RFC then supersedes the older one, and this is clearly noted on the front of the newer RFC. Another member of the IAB is the Internet Assigned Numbers Authority (IANA). The IANA is responsible for managing the list of values which make up the object identifiers used in the Internet protocol suite. For example, the IANA has assigned the number 1 to the RFC which defines the Internet standard MIB. Thus the object identifier for this RFC is mgmt(1), or 1.3.6.1.2.1.

A.1.1.1 How to Obtain a Copy of an RFC

The RFCs can be obtained through any of the following channels:

- Printed copies are available for a modest fee from the DDN Network Information Center:

```
Postal:  DDN Network Information Center
           142000 Park Meadow Drive
           Suite 200
           Chantally, VA 22021
           US

 Phone:  1 800-365-3642
             1 703-802-4535

  Mail:  nic@nic.ddn.mil
```

- In electronic form, users may use anonymous FTP (password: guest) to the host nic.ddn.mil (residing at 192.11.36.5) and retrieve files from the RFC directory.
- If your site doesn't have IP connectivity to the Internet community, but does have electronic mail access, an electronic mail message can be sent to the electronic mail address:

  ```
  mail=server@nisc.sri.com
  ```

 and in the subject field indicate the RFC number, for example, Subject: SEND rfcs/rfc1130.txt.
- If you have access to the World Wide Web, the RFCs can be obtained from the following address:

  ```
  http://info.internet.isi.edu
  ```

A.1.2 Functions of the IAB

As the coordinating committee for the Internet system, the IAB performs the following functions:

- Sets Internet standards
- Manages the RFC publication process
- Reviews the operation of the IETF and IRTF
- Performs strategic planning for the Internet, identifying long-range problems and opportunities
- Acts as a technical policy liaison and representative for the Internet community
- Resolves technical issues that cannot be treated within the IETF or IRTF frameworks

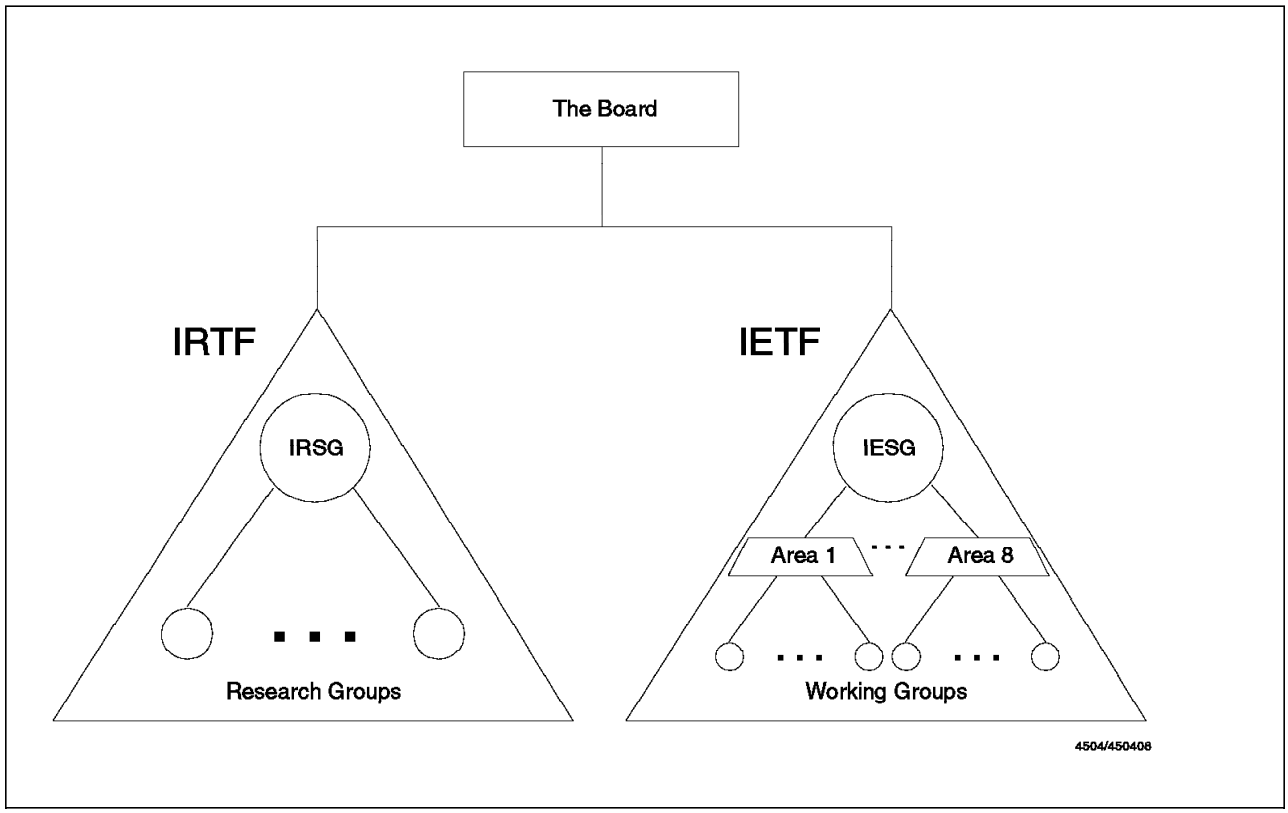


Figure 267. The IAB Organization

A.1.3 Protocol Standardization Process

The IAB provides standards with the intention of coordinating the evolution of the Internet protocols. With the increasing use of the Internet protocols for commercial purposes, standards coordination has become even more important.

Protocols that are to become standards in the Internet go through a series of states involving increasing amounts of scrutiny and experimental testing. At each step, the IESG of the IETF must make a recommendation for advancement of the protocol and the IAB must ratify it. This process is referred to as the standards track. If a recommendation is not ratified, the protocol is submitted

again to the IETF for further work. Table 52 on page 562 lists the Internet protocol state definitions.

Table 52. Internet Protocol State Definitions

Protocol State	Definition
Standard Protocol	The IAB has established this as an official standard protocol for the Internet.
Draft Standard Protocol	The IAB is actively considering this protocol as a possible standard protocol.
Proposed Standard Protocol	These are protocol proposals that may be considered by the IAB for standardization in the future.
Experimental Protocol	Typically, experimental protocols are those that are developed as part of an ongoing research project not related to an operational service offering. An experimental protocol may sometimes mean that the protocol is not intended for operational use. A system should not implement an experimental protocol unless it is participating in the experiment and has coordinated its use of the protocol with the developer of the protocol.
Historic Protocol	These are protocols that are unlikely to ever become standards in the Internet either because they have been superseded by later developments or due to lack of interest.

Table 53 lists the Internet protocol status definitions mentioned in A.1.1, “Request for Comments (RFC)” on page 560.

Table 53 (Page 1 of 2). Internet Protocol Status Definitions

Protocol Status	Definition
Required Protocol	A system must implement the required protocols.
Recommended Protocol	A system should implement the recommended protocols.
Elective Protocol	A system may or may not implement an elective protocol. The general notion is that if you are going to do something like this, you must do exactly this. There may be several elective protocols in a general area. For example, there are several electronic mail protocols, and several routing protocols.
Limited Use Protocol	These protocols are for use in limited circumstances. This may be because of their experimental state, specialized nature, limited functionality, or historic state.

Table 53 (Page 2 of 2). Internet Protocol Status Definitions

Protocol Status	Definition
Not Recommended Protocol	These protocols are not recommended for general use. This may be because of their limited functionality, specialized nature, or experimental or historic state.

Appendix B. A Brief Description of IBM Network Management Products

In this chapter we give a brief description of the products listed in 9.3, “Overview of IBM Products for Network Management” on page 402.

B.1 AIX Platform

These products work together to provide a management platform that is well suited for distributed and LAN Workgroup environments, and for interoperation with the MVS and OS/2 management platforms. A brief summary of the role of the various AIX systems management products is shown below:

- NetView for AIX

 NetView for AIX is an SNMP platform for managing heterogeneous, multi-vendor resources in distributed environments. It manages IP networks, SNMP devices and other non-IP resources. NetView for AIX provides configuration, problem and performance management functions, and an easy-to-use graphical user interface. It is an open platform with several interfaces for application integration, some of which are the End User Interface API, the Open Topology API, the SNMP API and the XMP API. Some of its highlights are:

 - Displays the network topology and monitors the status of devices
 - Integrates networks other than TCP/IP on its topology maps (for example, token-ring LANs) using its General Topology Manager component
 - Offers a consistent, graphical user interface for enhanced integration among applications and improved operator productivity
 - Supports non-IP environments using SNMP proxy agents (for example, NetBIOS PCs, token-ring LANs)

- LAN Management Utilities/6000

 LAN Management Utilities/6000 monitors and controls DOS, Windows and OS/2 systems, including OS/2 LAN Server and Novell NetWare servers. It uses LAN NetView Management Utilities for OS/2 as an SNMP proxy agent to support these PC environments.

- SNA Manager/6000

 SNA Manager/6000 interfaces with NetView for MVS to monitor and control an SNA subarea network from NetView for AIX. It displays graphical maps of the SNA physical and logical units.

- AIX NetView Service Point

 AIX NetView Service Point provides two-way connectivity to NetView for MVS. NetView for AIX uses NetView Service Point to send events as SNA alerts to NetView for MVS, and to receive commands from NetView for MVS.

- Router and Bridge Manager/6000

 The Router and Bridge Manager/6000 supports performance and error analysis for IP routers, including IBM 6611, Cisco and Wellfleet. It breaks down statistics by protocol and interface. It uses graphical displays and color-coded highlights on threshold exceptions.

© Copyright IBM Corp. 1996

- Nways BroadBand Switch Manager for AIX

 The IBM Nways BroadBand Switch Manager for AIX consists of several applications that gather, consolidate and process management data from BroadBand Network Switches. It supports fault, configuration, accounting, performance and operations management.

- LAN Network Manager for AIX

 LAN Network Manager for AIX monitors and controls token-ring LANs. LNM for AIX shows network topology and error information, and allows you to configure 8230 Controlled Access Units from NetView for AIX.

- Nways Campus Manager LAN for AIX

 IBM Nways Campus Manager LAN is a full management system for IBM 8250 Multiprotocol Intelligent Hub and IBM 8260 Multiprotocol Intelligent Switching Hub. It provides a rich set of management functions that enable the network administrator to have continuous visibility of the hub and its integrated concentration module status. It also is able to manage IBM 6611 Network Processor, IBM 2210 Nways Multiprotocol Router, and selected OEM routers. This Campus version includes APPN and Data Link Switch features.

 It further enables graphical remote control and monitoring of network devices including:

 – IBM 8224 Ethernet Stackable Hub
 – IBM 8230 (Models 3/13, 213, 4A/4P0 Token-Ring Concentrator
 – IBM 8271 (Models 001/108) EtherStreamer Switch
 – IBM 8272 (Models 108/216) LANStreamer Switch
 – IBM 8238 Nways Token-Ring Stackable Hub

- Nways Campus Manager ATM for AIX

 IBM Nways Campus Manager ATM for AIX is a state -of-the-art network management application package designed to manage your campus ATM network. It manages your ATM devices including:

 – IBM 8181 ATM LAN Bridge
 – IBM 8282 ATM Workgroup Concentrator
 – IBM 8285 Nways ATM WorkGroup Switch
 – IBM 8260 Multiprotocol Intelligent Switching Hub (when IBM Campus Manager LAN is additionally installed)

- LAN Remote Monitor for AIX

 The LAN Remote Monitor for AIX program offers a standards-based client/server solution that fits flexibly into your network. LAN Remote Monitor for AIX can collect data for any RMON-compliant management application and can direct any RMON-compliant probe.

 LAN Remote Monitor for AIX provides the following generic functions:

 – Full RMON support for token-ring and Ethernet LANs.
 – Summary screen gives you a high-level view of the entire LAN segment or ring.
 – Rapid fault discovery and response for identifying and solving network faults.

 - Graphical software for analyzing data and packets collected by remote probes.

- Nways Campus Manager for AIX

 IBM Nways Campus Manager Suite is a powerhouse suite that combines all the applications needed to manage traditional and ATM campus networks, and includes a remote networking monitoring (RMON) application. This suite contains:

 - IBM Nways Campus Manager LAN
 - IBM Nways Campus Manager ATM
 - IBM Nways Campus Manager Remote Monitor Advance

- Trouble Ticket for AIX

 Trouble Ticket for AIX is an application for recording, assigning and tracking problems. It has comprehensive inventory database and reporting capabilities. You can integrate Trouble Ticket with NetView for AIX to automate the creation of problem incidents from NetView for AIX traps.

- Systems Monitor for AIX

 The Systems Monitor product consists of two features:

 - System Information Agent - This is a smart agent that implements detailed instrumentation to monitor system processes and many system resources and attributes.
 - mid-level manager - It provides the ability to off-load polling activity from NetView for AIX to distributed AIX systems running this feature. The distributed Mid-Level Managers discover, poll and check thresholds for the devices on the local network segments, reducing network traffic and the load on NetView for AIX. IBM offers Systems Monitor for HP, Sun and NCR, in addition to AIX.

- AIX Transmission Network Manager/6000

 AIX Transmission Network Manager/6000 manages IDNX networks. It is a comprehensive network management solution that supports the growth of your corporate backbone network. TNM/6000 allows you to add new nodes without network management interruption, and provides real-time feedback on the status and configuration of the physical network.

- TMN Workbench for AIX

 The Telecommunications Management Network Workbench for AIX is a set of AIX-based tools for developing element, network and service management applications that use the OSI agent/manager model. It provides tools for the development of management applications and agents.

- NetView TMN Support Facility for AIX

 The NetView TMN Support Facility for AIX is an extension of NetView for AIX. It consists of the sum of the NetView for AIX functions and features added for the telecommunications industry. The TMN Support Facility includes the OSI stack support and a set of applications and services.

B.2 MVS Platform

A brief summary of the role of the various relevant products is shown below.

B.2.1 Basic Products

The following are the basic products on the MVS platform for managing heterogeneous environments.

B.2.1.1 NetView

The following are NetView for MVS components:

- Browse Facility
 - Allows you to look at NetView logs, data sets and files. Displays color-coded or highlighted messages in the log.
- Command Facility
 - Provides support for command processing, both those issued by an operator or issued by command procedures.
- Hardware Monitor
 - Collects, stores, and presents hardware and software statistical data, alerts and other error records. Forwards data to focal point NetViews.
- Session Monitor
 - Collects, stores and presents VTAM-owned resource data. Forwards data to focal point NetViews. Performs automatic node reactivation.

B.2.1.2 NGMF

- NGMF provides a graphic user interface for NetView based on GraphicsView/2, running on the OS/2 platform using Communications Manager/2.
- NGMF permits an operator to graphically view the status of resources in SNA and non-SNA networks and control them via generic commands.

B.2.1.3 RODM

- RODM is an in-memory data cache that provides a centralized source of information needed for systems management.
- Using an object-oriented approach, information is organized as objects in a hierarchical tree structure.
- The objects are categorized into different classes to provide inheritable characteristics.
- Alerts are generated by service point products running on AIX NetView Service Point, NetView/PC, or by applications like NPM or AOC/MVS, using the NetView program-to-program interface (PPI).
- RODM interfaces are provided to enable other management applications developed by IBM, vendors, or the customer to utilize the platform.

B.2.1.4 GMFHS

- NetView and GMFHS provide a graphic facility for managing and monitoring physical and logical resources of non-SNA resources.
- GMFHS extends the graphic capabilities of NGMF, providing the ability to control SNA and non-SNA resources.
- Integrated functions allow you to run commands against SNA and non-SNA resources from the graphical display using Command Tree/2, which provides all the commands that can be used for specific resources.

B.2.1.5 MSM

- MSM enables centralized management of LAN Management Utility (LMU) based networks, LAN Network Manager (LNM) and Novell NetWare managed networks, and IP networks from the NGMF workstation.
- MSM provides dynamic topology, status discovery and appropriate command sets to manage and monitor IP, LMU, LNM and NetWare networks. By providing these functions, MSM simplifies managing the entire network resources.
- IBM NetView MultiSystems Manager is available as a base feature, and additional topology features for IP, LMU, LNM and NetWare resources.
- With MSM you can manage the following network resources:
 - IP networks
 - IP segments
 - IP routers
 - IP hosts
 - IP hubs
 - LMU managed LAN servers
 - LMU managed NetWare server
 - TR LAN adapters
 - TR LAN bridges
 - TR LAN segments
 - TR LAN controlled access units (CAUs)
 - NetWare servers

B.2.2 Optional Products

The following are optional products to manage from the MVS platform.

B.2.2.1 AIX NetView Service Point and NetView/PC

AIX NetView Service Point allows AIX environments to exchange network management information with NetView.

NetView/PC is based on an OS/2 workstation.

These service points collect network management information and forward it to NetView for centralized management, and provide the option of using a service point for controlling non-SNA devices.

NetView/PC provides API/CS programming interface to write applications and alert logging.

Many vendors provide applications running on NetView/PC or AIX NetView Service Point.

B.2.2.2 Six2View

Six2View integrates the control of a DEC network environment into NetView providing functions to monitor, control and manage DEC resources.

Six2View is tightly integrated with NetView and provides:

- Direct interface between NetView and DECmcc
- NetView operator command and response panel
- NetView automation
- Generate alerts by DECmcc and send to NetView
- Works together with DEC VMS (commands from NetView to any VMS)

Six2View using GMFHS and RODM provides a graphical display on NGMF and uses Command Tree/2.

B.2.2.3 NetView Bridge

- The NetView Bridge is a standard part of NetView and the Information Management program.
- NetView Bridge provides two-way access between the Information/Management database and NetView, for problem and configuration purposes.
- The operator using NGMF can access information/management problem records for specific devices to display, update or open a problem with IPM. In addition configuration data can be accessed.

B.2.2.4 NetView AutoBridge/MVS

- NetView AutoBridge/MVS is based on the NetView Bridge and automates the process to open problem records in the information/management database from NetView.
- Autobridge links together NetView and information/management.
- Problem records can be opened automatically when an alert or error message is received in NetView.

B.2.2.5 NetView APPNTAM

- NetView APPNTAM provides functions for managing APPN network environments.
- APPNTAM collects and stores APPN topology data in RODM, including real-time updates in the RODM data cache.
- APPNTAM provides a dynamic, graphical display of APPN topology, using NGMF.
- Control of the SNA ports and links is provided by using commands on the pull-down menus or Command Tree/2 at the NGMF workstation, or by using native commands on operator console.

- APPNTAM allows centralized collection on APPC sessions and conversation accounting information.

B.2.2.6 AON/MVS

- AON/MVS is based on NetView provides automated network operations management system for the MVS environment.
- AON/MVS provides common routines such as generic failure routine, generic recovery routine and generic active recovery routines.
- The automated recovery of network resources is based upon VTAM messages, management service units (MSUs) and monitoring activities.
- AON/MVS comes with the AON/MVS base feature and additional components for SNA, LAN and TCP/IP automation.
- AON/MVS provides a single point of control, based on a 3270 display interface, pro-active help desk facilities, and the reporting facility.
- The AON/SNA feature intercepts over 40 critical VTAM messages and alerts for network resources. AON/SNA issues commands against the failing resources to reactivate and monitor them until they are active again. AON/SNA automates SNA, Switched Network Backup (SNBU) and X.25 resources. The actual status of the resources is displayed on the Dynamic Display Facility (DDF) and can be displayed in exception mode only.
- The AON/LAN feature monitors token-ring networks communicating with IBM LAN Manager, or IBM LAN Network Manager. The following are some examples of AON/LAN functions:
 - Shows status changes in the LAN environment
 - Automatic recovery of bridge links connecting token-ring segments
 - Communicates and automates commands to LAN Network Manager
 - Communicates and automates commands to LAN Server Program
 - Communicates and automates commands to LAN Management Utility
 - Communicates and automates commands to Remote Operations Service
 - Communicates and automates commands to IBM Bridge Program
- The AON/TCP feature monitors TCP/IP networks communicating with NetView for AIX. NetView for AIX sends alerts to notify NetView for MVS of TCP/IP resource status changes. AON/TCP detects and reacts to TCP/IP resource failures.
- AON/TCP checks for performance problems, such as CPU utilization and disk space utilization, name server failure in the TCP/IP network, unavailable resources and security authorization failures.

B.2.2.7 NPM

NetView Performance Monitor belongs to the NetView family. Using NPM you can perform the following tasks:

- Collect and monitor performance data
- Collect accounting data
- Determine problems in the network

NPM provides extensive statistics on network traffic, queue length, buffer utilization, communication controller activity and VTAM buffers and data. In

addition NPM collects data from local area networks (LAN) and Novell NetWare resources, Response Time Monitor (RTM) data, X.25 traffic data, frame relay data, NEO resource data and response times.

The collected information data can be displayed in table format using 3270 format or graphic format using GDDM, or the Desk/2 interface, running on OS/2.

NPM Desk/2 running as an APPC program on an OS/2 workstation provides functions to collect and display the performance data using Configuration and DataView windows based on the presentation manager technology of OS/2 2.1.

NPM can collect end-to-end response times on sessions through any session manager, including TPX.

The NPM Batch Reporting Facility allows you to run reports against collected information.

B.2.2.8 NetView Remote Operations Manager MVS and NetView Remote Operations Agent/400

NetView Remote Operations for AS/400 provides centralized management of AS/400 systems from the MVS platform.

Remote Operations for AS/400 uses the architected OPERATE command and provides a broadcast capability to send commands to all AS/400 systems.

The commands to the AS/400 can be sent in two ways:

- Using the automation facilities from NetView
- Using the NetView command line

Using the NetView Remote Operations Agent/400, the AS/400 system catches all commands coming from NetView and passes status information and responses back to NetView for MVS.

B.2.2.9 NetView Network Planner/2 (NNP/2)

NetView Network Planner/2 is an OS/2-based product managing enterprise-wide inventory and assets of networks, systems, and the resources within them. NNP/2 includes the following functions:

- Data model support for software, equipment, features, locations, circuits, organizations, people and financial information using SQL.
- Facility to import/export data from/to RODM and supply data to Trouble Ticket/6000.
- Configuration data provided by LMU and LNM can be stored within NNP/2.
- Data can physically be stored in any DB2 or DRDA supported database.
- Easy-to-use tool for displaying and changing planning information.
- Graphical support for easily visualizing complex logical and physical relationships.

B.2.2.10 Teleprocessing Network Simulator (TPNS)

TPNS allows you to simulate application and network traffic, to examine system and network performance.

B.3 OS/2 Platform

These products work together to provide a management platform that is well suited for LAN Workgroup environments, and for interoperation with the AIX and MVS management platforms. A brief summary of the role of the various OS/2 systems management products is shown below:

- NetView for OS/2:
 - Delivers a multi-vendor, open systems and network management platform for client/server environments.
 - Implements the industry-standard SNMP protocol. It can manage any device with an SNMP agent.
 - Provides SNMP APIs to allow vendors and customers to write LAN systems management applications.
 - Supports SNMP natively over TCP/IP, IPX and NetBIOS, and over SNA with AnyNet/2.
 - Offers a consistent, graphical user interface for enhanced integration among applications and improved administrator productivity.
 - Provides SNMP applications (fault, configuration and performance) and tools (MIB loader, MIB browser, MIB application builder and MIB data collector).
 - LAN NetView Management Utilities is packaged with NetView for OS/2.
 - Interoperates with NetView for MVS and NetView for AIX.
- LAN NetView Management Utilities (LMU):
 - Provides utilities for managing PCs in NetBIOS and IPX LANs.
 - Offers functions for operations, problem, performance and configuration management.
 - Includes an SNMP proxy agent to communicate with SNMP managers, such as NetView for AIX and NetView for OS/2.
 - Interoperates with NetView for MVS and the Multisystem Manager.
- IBM SystemView for OS/2:
 - Offers tools for managing PCs in TCP/IP, NetBIOS and IPX LANs.
 - Emphasizes ease of use and hardware management.
 - Monitors PC hardware components.
 - Provides hardware and software configuration information.
- LAN Network Manager for OS/2 (LNM):
 - Manages token-ring LAN media (network adapters, bridges and hubs), with some support for Ethernet LANs using bridges.
 - Provides a graphical view of the LAN topology.
 - Includes a proxy agent to communicate with LAN Network Manager for AIX (which runs on NetView for AIX).

 - Interoperates with NetView for MVS and the Multisystem Manager.
- DatagLANce:
 - Captures, monitors and analyzes Ethernet and token-ring LAN data.
 - Decodes multiple network protocols in real time.
 - Sets filters and event detectors for data capture.
 - Records network traffic data and plays it back later for analysis.
 - Loads the LAN for stress testing.
- System Performance Monitor/2 (SPM/2):
 - Measures, collects and reports OS/2 performance data.
 - LMU works with SPM/2 to monitor OS/2 performance thresholds.
- Distributed Console Access Facility (DCAF):
 - Allows complete control of remote PCs. The DCAF operator can take over the keyboard, display and mouse of another PC.
 - Supports TCP/IP, IPX, NetBIOS, SNA and ASYNC connections.
- NetView Distribution Manager/2 (NVDM/2):
 - Offers software distribution and change control functions for PCs in NetBIOS LANs.
 - Supports both CID (Configuration Installation Distribution) and non-CID (replication) software installations.
 - Interoperates with NetView Distribution Manager for MVS.
- NetView DM Easy Preparer for OS/2 (EasyPrep):
 - Simplifies the preparation of software packages to be distributed and installed on OS/2, DOS and Windows systems using NVDM/2.
 - Automates the generation of response files (CID-enabled software), modification files (non CID-enabled software) and change files before being distributed with NVDM/2.
 - Saves CID-enabled software configuration data in a relational database that can be reused for future installations.
 - Collects software and node definitions directly from the NetView DM for MVS SPMF (Software Profile Management Facility) DB2 database.
- Network Door/2 (NetDoor):
 - Delivers an application serving infrastructure for OS/2, DOS and Windows applications in TCP/IP and NetBIOS LANs.
 - Makes applications easily available to the end users via a catalog, without installing them on their PCs. The applications are run from a NetDoor server.
 - Implements centralized maintenance and administration, providing a single system image for LAN applications.
- ADSTAR Distributed Storage Manager/2 (ADSM/2):
 - Provides backup, archive and restore functions to protect data stored on PCs and UNIX systems.
 - Supports TCP/IP, SNA, NetBIOS and IPX networks.

- Delivers automated, policy-based data management.

B.4 Windows Platform

These products provide a management platform that is well-suited for small LAN Workgroup environments, but they do not interoperate with the other IBM management platforms. A brief summary of the role of the IBM Windows-based systems management products is shown below:

- NetView for Windows:
 - Delivers a multi-vendor, open network management platform for LAN interconnected environments.
 - Supports the industry-standard SNMP protocol. It can monitor any device with an SNMP agent.
 - Focuses on managing network devices, not PC systems.
 - Integrates PSMs (Product Specific Modules) and PIMs (Product Integrator Modules) that manage specific network devices from multiple vendors. These PSMs and PIMs provide two levels of function:
 - Product Integrator Modules (PIMs)

 PIMs interact with the operator using menu panels. They may optionally offer a picture of the device without any support to interact with the operator.
 - Product Specific Modules (PSMs)

 PSMs provide more advanced management functions and interact with the operator using both menu panels and a graphical picture of the device. This graphical representation shows status information and may include the actual control panel, interfaces, ports and switches of the device.
 - Stores its fault, performance and configuration data in an object-oriented database.
 - Offers a consistent, graphical user interface for enhanced integration among device management applications and improved administrator productivity.
- NetFinity Manager for Windows:
 - Offers tools for managing PCs in TCP/IP, NetBIOS and IPX LANs.
 - Focuses on managing PC systems, not network devices.
 - Emphasizes ease of use and hardware management.
 - Monitors PC hardware components.
 - Provides hardware and software configuration information.
- LAN Remote Monitor for Windows:

 The LAN Remote Monitor for Windows program offers a standards-based client/server solution that fits flexibly into your network. LAN Remote Monitor for Windows can collect data for any RMON-compliant management application and can direct any RMON-compliant probe.

 LAN Remote Monitor for Windows provides the following generic functions:
 - Full RMON support for token-ring and Ethernet LANs

 - Summary screen gives you a high-level view of the entire LAN segment or ring
 - Rapid fault discovery and response for identifying and solving network faults
 - Graphical software for analyzing data and packets collected by remote probes

- Nways Manager for Windows

 The Nways Manager for Windows will help you manage your campus network environment easier and more effectively than ever. It is an integrated suite of network management applications that work seamlessly with IBM NetView for Windows management platform to remotely control and monitor networking devices such as:

 - IBM 8224 Ethernet Stackable Hub
 - IBM 8230 (Models 3/13, 213, 4A/4P) Token-Ring Concentrator
 - IBM 8238 Nways Token-Ring Stackable Hub
 - IBM 8271 (Model 001/108) EtherStreamer Switch
 - IBM Turboways 8282 ATM Workgroup Concentrator
 - IBM 8281 ATM LAN Bridge
 - IBM 8250 Multiprotocol Intelligent Hub
 - IBM 8260 Multiprotocol Intelligent Switching Hub
 - IBM 6611 Network Processor
 - IBM 2210 Nways Multiprotocol Router

Appendix C. IBM infoMarket Rights Management Architecture

From an architectural perspective, the IBM infoMarket service will use multiple technologies to implement the rights management system. These include:

Access Control: IBM infoMarket will perform an access permissions evaluation to determine user rights to access information sources and documents. This access evaluation will be performed based on user credentials and the user profile maintained in the IBM infoMarket service registry.

Authentication: IBM infoMarket validates a consumer's right to use the service through an authentication process. Authentication in many Web applications (including the IBM infoMarket service) is performed today through the use of a user ID/password assigned to a consumer during service registration. Future industry direction for authentication involves the use of digital IDs, or certificates. A certificate is the equivalent of your digital driver's license; it validates that you are who you say you are and your affiliations and restrictions.

Additionally, an information provider document server that wants to communicate with an IBM infoMarket server will be authenticated before access is granted.

Cryptolope containers will be authenticated through digital signatures to ensure that the consumer is receiving the original unaltered document that was requested.

Browsers: The IBM infoMarket service can be accessed through Web browsers such as Netscape Navigator or Mosaic or through custom applications created using the IBM infoMarket Client Toolkit.

Web browsers access the IBM infoMarket service through the IBM infoMarket servers on the World Wide Web. Secure transactions will take place between the browser and the servers when the consumer is using a secure″Web browser such as Netscape Navigator V1.22 (which supports the Secure Sockets Layer (SSL) protocol). IBM infoMarket plans to provide a browser helper application to handle the document buy interaction with the consumer, and then return the decrypted document from the Cryptolope container to the browser. IBM infoMarket will also work with leading browser developers to add native support for Cryptolope containers.

Custom IBM infoMarket applications will access IBM infoMarket servers using an ″under-the-covers″ message protocol within the Client Toolkit. These custom applications will allow an end-user interface targeted to a specific customer audience to access specific content. While today's browsers allow a consumer to print and save any document being viewed in the browser, a custom application can control the actions the consumer is allowed to perform based on permissions and usage fees the customer has agreed to pay.

Clearing Center: The IBM infoMarket clearing center is responsible for evaluating rules, defined by the content provider, to determine the consumer's permission to access and use information content. The clearing center will also generate activity records for reporting and accounting.

© Copyright IBM Corp. 1996

Content Rating: Content rating is the filtering of content based on specific attributes assigned to information. The electronic commerce industry is currently looking to establish standards for labeling content. Once these standards are written, browsers and custom applications may be developed to enforce them.

Cryptolope Containers: See Secure Containers.

Encryption: The IBM infoMarket service will use standard encryption technologies to prevent unauthorized access to a document, including cryptography and digital signatures. See Secure Containers for more information.

Event Management System: The IBM infoMarket event management system is responsible for logging activities such as use of content. The event logs that it generates are subsequently used for accounting and reporting activities.

Fingerprinting: A digital fingerprint is an invisible record of who ″touched″ an electronic document. When implemented, fingerprinting will allow the IBM infoMarket service to determine who first misused the document (such as by improperly copying or distributing it). Fingerprinting electronic documents is content-dependent and is easiest for document types that have a large number of bits. Only a small number of bits are needed to identify the ″culprit;″ they are not noticeable in the background of a picture and are hidden in the least significant bits of digital audio or video. Low-bandwidth data, such as ASCII text files, are more difficult to fingerprint. Fingerprinting will be selectable by the content provider.

Metering: Using the IBM infoMarket event management system, applications can record content usage of Cryptolope containers and associated actions.

Offline Support: Over time, IBM infoMarket will use secure payment cards, smart cards, and other state-of-the-art technologies to support the disconnected user.

Rights Management Language: The IBM infoMarket rights management language specifies the rules for determining the consumer actions permitted for accessing a document and the costs associated with consumer actions. In addition, the rights management language will be used to specify the consumer′s permission to access a particular information source and document.

Secure Containers (Cryptolope Containers): The IBM infoMarket service will use a secure container architecture to package and distribute information content and properties. We call this container a cryptographic envelope, or Cryptolope container. A Cryptolope container holds an encrypted versi on of a document (a document may contain many data formats such as ASCII text, HTML, image) and information for evaluating the consumption of a document such as an abstract, actions that can be performed on it, the associated costs and copyright notices. Once the content is purchased, the IBM infoMarket service will transparently provide the customer with a private key to unlock the Cryptolope container.

Cryptolope containers can be issued and handled by multiple issuing authorities. This will ensure an open standard for Cryptolope

containers and will allow a free market for implementing custom client applications.

Cryptolope containers may be disseminated widely by using alternative distribution methods. Because Cryptolope containers can be large, they can be distributed using an inexpensive but insecure method. On the other hand, the keys are quite small, so an expensive but secure distribution method can be used for them. When the user's browser points to an IBM infoMarket Cryptolope container, helper applications talk to the IBM infoMarket client code that asks for the key to be delivered over the Internet using session security. The content is opened and decrypted in the client machine, and then delivered to the appropriate viewer.

Because the Cryptolope container is a document in digital form, it will use digital authentication techniques. Digital signature for authentication is a standard technique in the public key encryption repertoire.

Secure Hardware Environment: Secure hardware will be used as an alternative form of payment, either digital "cash" or digital credit. A device such as a smart card can be filled up with digital currency from some kind of digital bank. Then it can be used to pay for IBM infoMarket purchases and a record of content purchases can be maintained.

Transport Level Security: The IBM infoMarket service will use a transport layer security mechanism (for example, the Secure Sockets Layer (SSL) protocol) to prevent unauthorized access to important information such as credit card numbers. This mechanism uses cryptography so that information is not in the "clear" while being transmitted across a network.

Watermarking: In a paper document, a watermark is a physical design embossed or pressed into the paper that can be seen when the page is held up to a light. In an electronic document, it is usually a faint background image superimposed over the document image. In an electronic document, the main function of a watermark is to make visible on every page that the document is copyrighted.

Appendix D. More Information about IBM infoSage

The following contains additional information on infoSage:

D.1 Content Resources

IBM infoSage provides information from the following resources:

- COMTEX is a leading aggregator of hundreds of real-time news sources from around the world. COMTEX provides news and information from the following sources:
 - A & G Information Service: business, political and economic news from the former Soviet Union and Eastern Europe.
 - Africa News Service: news coverage from the African continent.
 - American Banker/Bond Buyer: news of the banking and bond markets.
 - AsiaInfo Services: news abstracts from over 600 local newspapers and journals across China.
 - Business Wire: full-text corporate press releases for 12,000 U.S. companies.
 - Cineman Syndicate: latest music, video and book reviews.
 - COMTEX Newsroom: news coverage of the major headlines of the day and up-to-the-minute financial news and statistics on domestic and foreign markets.
 - FedNet Government News: abstracts of ″The Congressional Record″ and ″The Federal Register″
 - Futures World News: news and information on commodities traded on the world′s commodity futures exchanges.
 - Inter Press Service: providing news originating in developing and third-world nations.
 - ITAR/TASS News Agency: news, business and sports from Russia.
 - Knight-Ridder/Tribune Business News: a leading domestic newswire service that provides timely business news from more than 70 newspapers and magazines throughout the U.S., providing selected items, on a daily basis, from Knight-Ridder Fina, an around-the-clock service that reports on business, finance and economic news.
 - Knight-Ridder/Tribune News Service: news, features, sports and financial coverage from some of America′s best newspapers along with a global perspective from correspondents based in Europe, the Orient, Middle East, Africa and Latin America.
 - Pan-African News Agency: news from across Africa covering 48 national news agencies.
 - PR Newswire: full-text corporate press releases from over 17,000 U.S. companies.
 - South American Business Information: providing daily news abstracts from Argentina, Brazil, Chile, Paraguay and Uruguay.

© Copyright IBM Corp. 1996

- The Sports Network: up-to-the-minute coverage of all domestic and international sports events.
- United Press International: up-to-the-minute news, business and sports stories from around the world.
- U.S. Newswire: full-text press releases from U.S. government agencies.
- Xinhua News Agency: news coverage from all 30 provinces in China as well as Hong Kong, Macao, Latin America, the Middle East and Africa.
- Ziff-Davis Wire Highlights: timely articles on the high-tech industry including the movers and shakers in the industry and important stories that affect members.

• Dun & Bradstreet Corporation is the world′s leading marketer of information, software and services for business decision making, with 1995 revenues of $5.4 billion.

 D&B will provide marketing information from Dun′s Market Identifier, which accesses 10.2 million company records from D&B′s United States database.

• Information Access Company: headquartered in Foster City, CA, is an indirect subsidiary of the Thomson Corporation. Information Access Company is acknowledged as one of the premier providers of electronic information to corporations through the world, delivering timely information-based solutions to several million people each day via online services, CD-ROM products, and magnetic tape for LAN/WAN delivery to the desktop. Members of IBM infoSage have access to abstracts and full text from the following databases:

 - IAC PROMT: International in scope and outlook, PROMT covers 65 major industries, offering substantive information about companies, the products and technologies they develop, and the markets in which they compete. PROMT is comprised of business journals, newsletters, and newspapers.
 - IAC Trade & Industry Database: Focuses on market and industry trends, management concerns and challenges, legislative and regulatory decisions, global economic conditions, and corporate profiles. They contain more than 1,000 sources include trade and business publications, regional business journals, and economic and management journals.
 - IAC Newsletter Database: A virtual library of more than 600 full text newsletters, offering expert opinions, analysis, and inside information on industries and business activities spanning five continents.
 - IAC Magazine Database: An online barometer of popular culture that provides current and retrospective news from hundreds of popular magazines and newsstand publications, focusing on consumer behavior and lifestyles, media trends, politic opinion, and leisure activities.
 - IAC Health Periodicals Database: Provides coverage of consumer health and professional medical journals, specifically in the areas of health, medicine, fitness and nutrition. It contains information from 100 leading consumer health and professional publications.
 - IAC Computer Database: Features a collection of over 100 leading business and consumer publications, and identifies product evaluations, trade names, user techniques, language revisions, and computer industry standards and specifications.

 - IAC Industry Express: Features a rolling 30 days content from leading trade and business publications in IAC PROMT, IAC Trade & Industry Database, IAC Newletter Database, and IAC Computer Database -- at or near the time of their original publication.

- Intell.X: A division of DataTimes, one of the world′s largest business information providers, is headquartered in Arlington, Virginia.

 Intell.X will provide information sources from SourceEXpress, which is comprised of full text business articles from over 300 news sources including newswires, regional, national and international newspapers, trade and business magazines, journals, and industry reports.

 Intell.X offers outstanding international and domestic sources and uniquely strong leading regional news sources, such as the Cincinnati Post, Chicago Sun-Times, Salt Lake Tribune, St. Petersburg Times, The Daily Oklahoman, Allentown Morning Call, and the American Cities Business Journals.

- PAWWS: A division of Security APL, is the first company to provide stock quotes and the ability to trade online via the Internet. PAWWS is an Internet-based company providing portfolio accounting, online trading, stock quotes, news, data and other financial information to individual investors.

 PAWWS provides IBM infoSage members with 20-minute delayed stock quotes for U.S. stocks traded on the major stock exchanges. Information on mutual funds and money markets will also be available. In addition, with Stock Tracker, subscribers can select up to 20 stocks and receive closing quotes from the day before in their morning delivery.

- The Reference Press: The nations leading provider of company information to consumers and professionals, provides access to Hoover′s Company Profile Database, which includes more than 1,800 in-depth profiles of leading public and private U.S. and global companies.

 Hoovers company information is the most affordable source for information on the operations, strategies, histories, financial performance and products of major U.S. and global public and private enterprises. This information is available in print, online, facsimile, personal digital assistants, CD-ROM and diskette formats, and through the Hoover′s Online site on the World Wide Web (http://www.hoovers.com).

- Reuters NewMedia Inc.: a U.S.-based subsidiary of Reuters Holdings PLC, one of the world′s largest news organizations, provides a wide range of news and news picture services:

 - The Reuters Online News Service: a package of top 10 news stories, updated hourly in five subject categories, U.S. news, international news, business, sports, and entertainment news.

 - The Reuters World Service: An English-language, international general news wire.

 - The Reuters North American News Report: A general news service featuring the day′s top North American stories reported by 21 Reuters bureaus in the U.S. and six in Canada.

 - The Reuters Business Report: American, European and Asian versions of a popular daily news wire which Reuters supplies to general and business media worldwide.

 - The Reuters/Variety Online Entertainment Report: A real-time news service edited in Hollywood covering all national and international aspects of the entertainment industry.
 - The Reuters Corporate World News Service: a daily international news service developed for corporate executives and business analysts, covering more than 12,000 companies.
 - The Reuters News Picture Service: a daily feed of news pictures from U.S. and international datelines featuring general, sports, and entertainment coverage.

- Standard & Poor′s: A division of The McGraw-Hill Companies, provides financial, economic, and investment information, as well as analytical services, to the global financial community and commodity trading markets.

 Standard and Poor′s provides IBM infoSage members with access to The Standard & Poor′s Register of Corporations, a listing of executive rosters, addresses, and telephone numbers from 55,000 public and privately-held companies. The register also contains principal products, SIC (industry) codes, number of employees, annual sales, and names of the primary accounting firm, bank and law firm of each company.

- Weather Services Corporation: Is one of the oldest and most respected sources of worldwide commercial weather information. Weather Services Corporation provides specialized and customized meteorological information to vertical markets including agribusiness, electric and gas utilities, state, municipal and local governments, newspapers, general public and industry specific news services, the broadcast industry, and marine industry.

 Weather Services Corporation will provide specialized information including national and regional forecast information, 5 day forecasts for major business locations, hurricane and tropical analysis, business travel forecasts, significant weather alerts, and breaking weather stories as they happen.

D.2 IBM Profile Editor Screens

The first panel you see is the business topics screen.

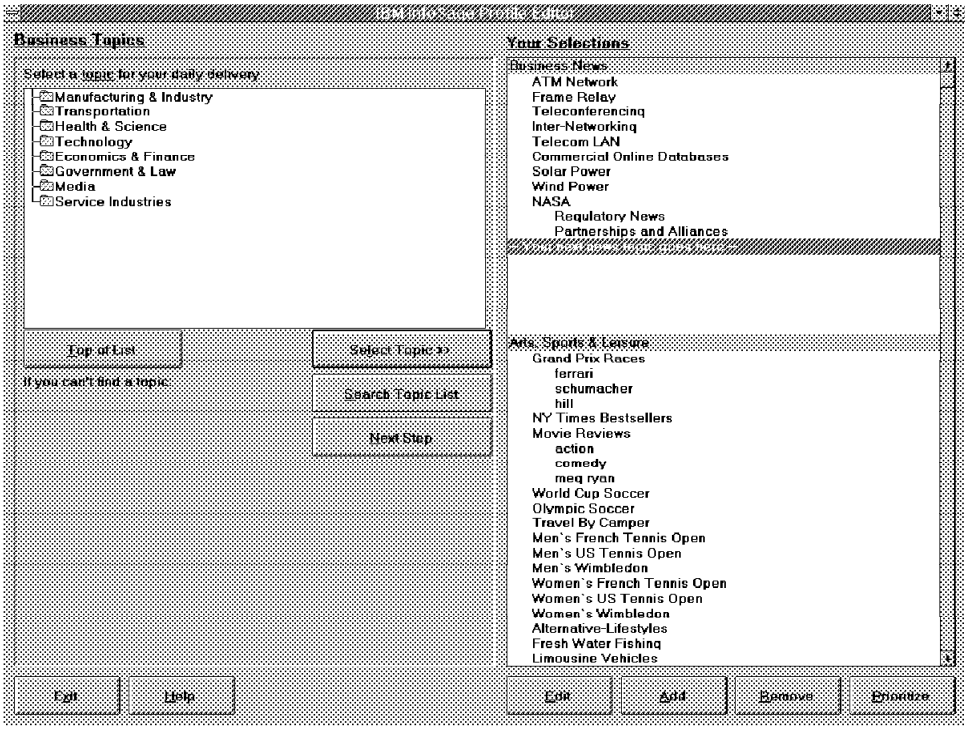


Figure 268. IBM Profile Editor Business Topics

When you select a topic you will come in the personalize window. Here you can enter three words or phrases to personalize the business topic.

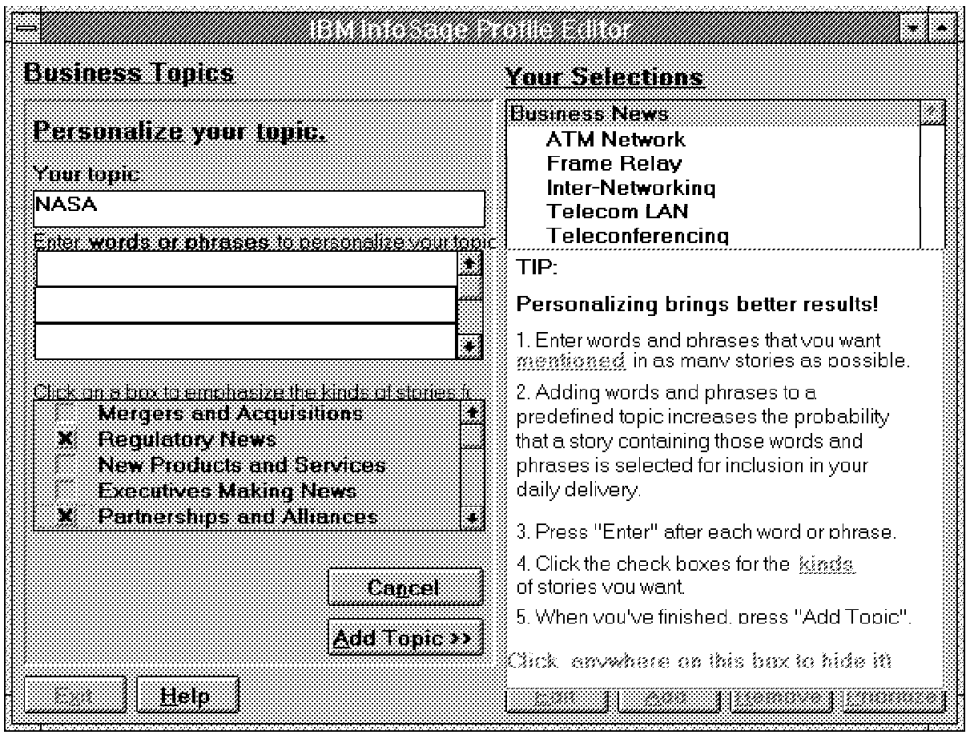


Figure 269. Personalized Business Topics

After the business topics you can select your leisure topics.

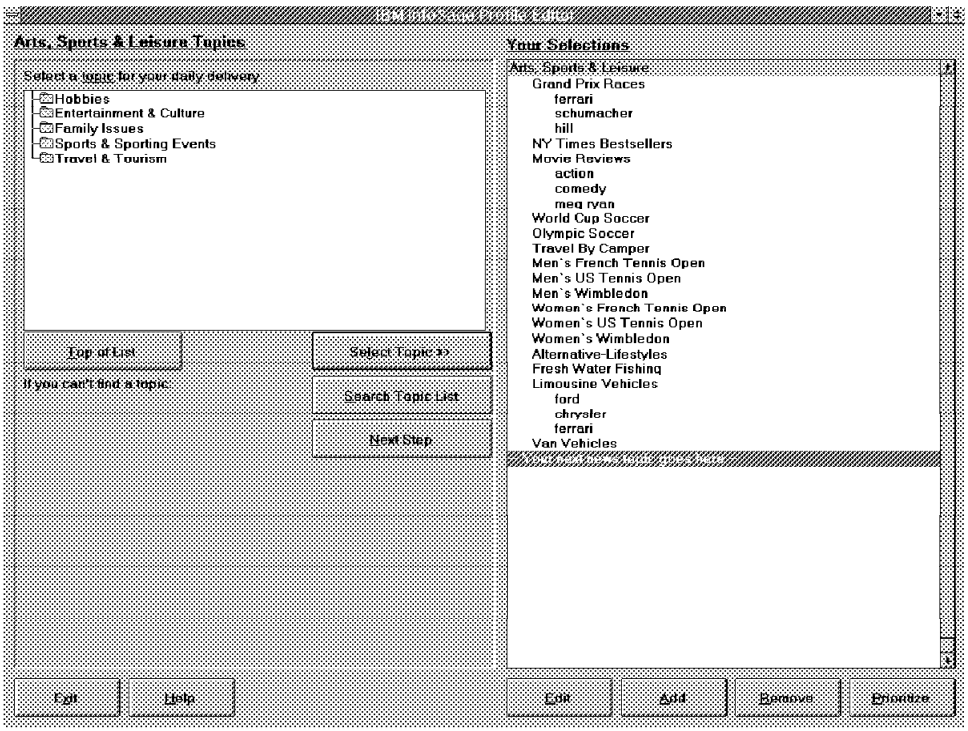


Figure 270. IBM Profile Editor Leisure Topics

At the end of the topic selection you can prioritize your topics and set alarm filters to different topics. In Figure 271 you can see this alarm settings in the first tw topics in the right box.

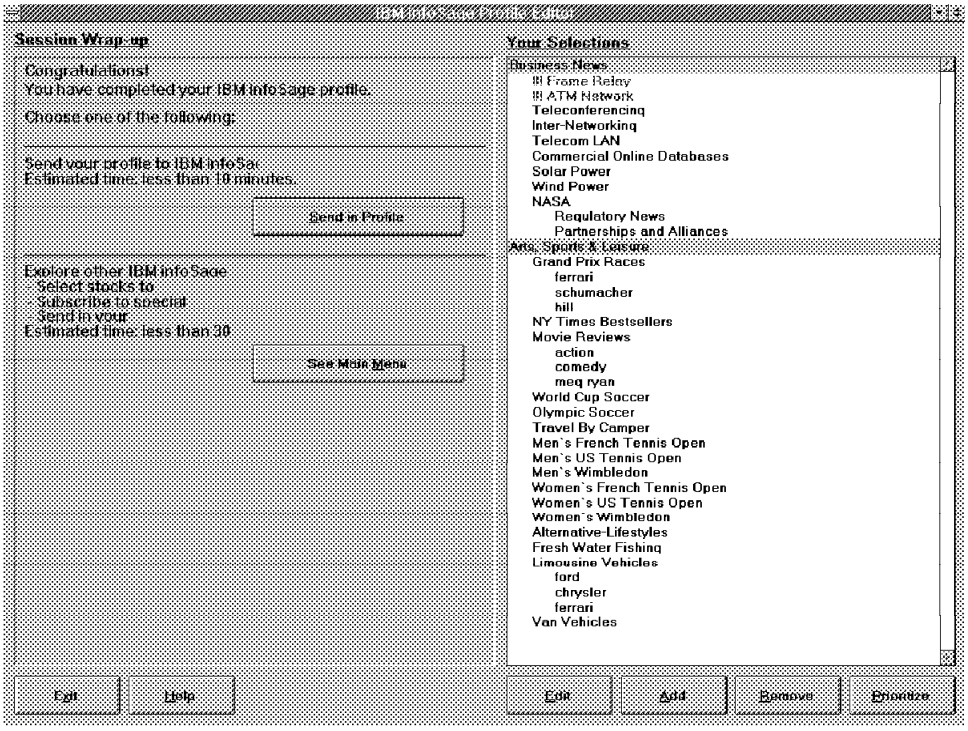


Figure 271. IBM Profile Editor Session Wrap-Up

Further possibilities of the IBM infoSage system are the stock tracker and the Special Editions functions. In the stock tracker you can select up to twenty different stocks and you will get daily information about them.

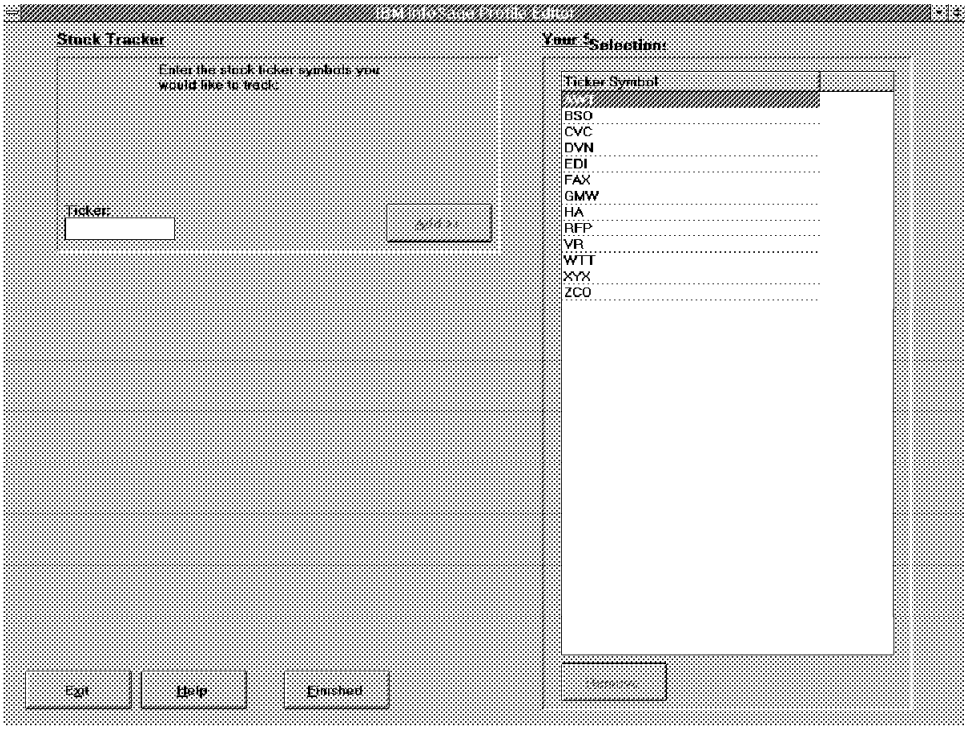


Figure 272. IBM Profile Editor Stock Tracker

In the Special Edition part you find daily, weekly and monthly editions in different areas of interest.

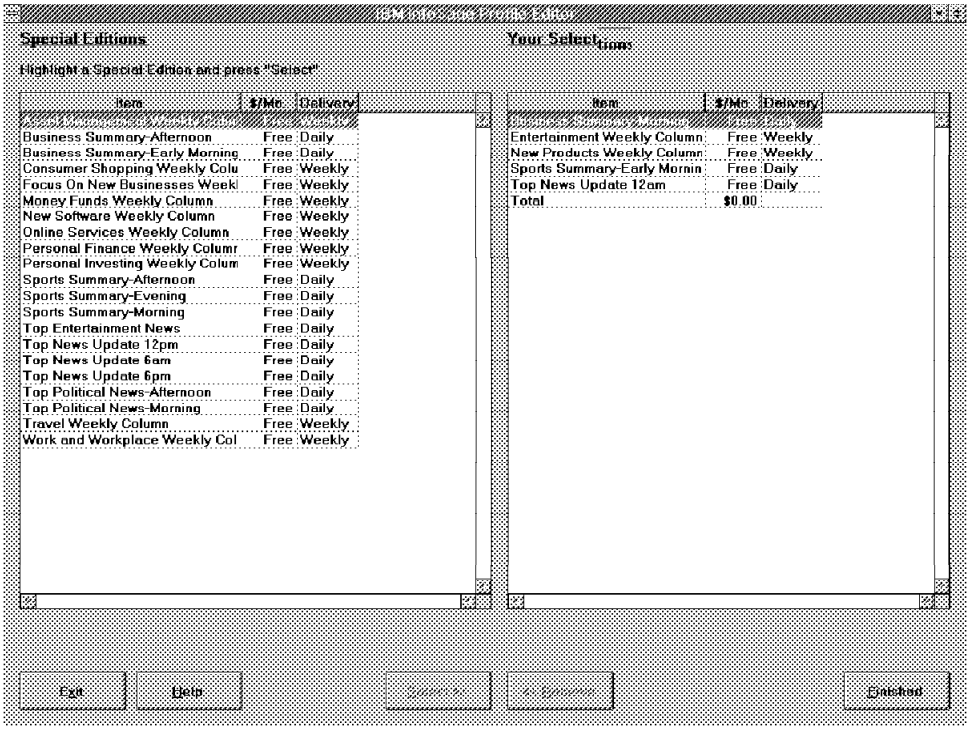


Figure 273. IBM Profile Editor Special Editions

D.3 IBM infoSage Result Examples

Here are three examples of information delivered from IBM infoSage:

D.3.1 Morning Delivery

```
Subject: IBM infoSage Morning Edition (June 27) <179CAAAABKXV>
Date: Thu, 27 Jun 1996 02:29:36 -0400
From: sage@chicory.infosage.ibm.com
To: DEIBMC5K@IBMMAIL.COM
Subject: IBM infoSage Morning Edition (June 27) <179CAAAABKXV>
**************************************************************************
                      IBM infoSage Morning Delivery
**************************************************************************
Patrick Schmitt-Heinrich
Member ID: XXXXXXXXX

Company: IBM XXXXXXXX
**************************************************************************
Thursday, 27 June 1996
02:32 AM EDT

Document ordering instructions and member services information are
located at the end of your delivery.

**************************************************************************
I. YOUR BUSINESS NEWS
**************************************************************************

1. GMP's Terreri becomes president of national wind power organization
   Topic Matched: Wind Power

 SOUTH BURLINGTON, VT. (June 26) BUSINESS WIRE -June 26, 1996--A. Norman
 Terreri, executive vice president and chief operating officer of Green
Mountain Power Corp., today became president of the American Wind Energy
 Association (AWEA), a national organization supporting the development
of wind power. Terreri began his one-year term at WINDPOWER '96,...

< >Full text of this story - FREE


2. 'Live from Mars' telecasts
   Topic Matched: NASA

PASADENA, June 26 (UPI) _ Educators and students can become virtual
travelers
to Mars|

< >Full text of this story - FREE

**************************************************************************
II. YOUR ARTS, SPORTS & LEISURE NEWS
**************************************************************************

1. Former President May be Abiola Suspect?
   Topic Matched: Alternative-Lifestyles

Lagos (by Godwin Agbroko), June 26, 1996 via Africa News - The
administration's suspicions turn to former   President Ibrahim Babangida
```

on the June 4 Murder of Kudirat Abiola. Is this a case of calling a dog a bad name?

< >Full text of this story - FREE

2. Seles Knocked out of Wimbledon Tennis Championships
 Topic Matched: Women's Wimbledon

WIMBLEDON, ENGLAND, June 26 (Xinhua/OANA) -- The world joint number one tennis player Monica Seles of the United States was defeated by unseeded Slovak player Katarina Studenikova in the second round of the on-going 1996 Wimbledon Tennis Championships.

< >Full text of this story - FREE

3. TENNIS: SELES STUNNED AT WIMBLEDON
 Topic Matched: Women's Wimbledon

The major upset bug made its way into the women's draw at Wimbledon on Wednesday when second-seeded and world co-number one Monica Seles was stunned by Slovakia's Katarina Studenikova, 7-5, 5-7, 6-4.

< >Full text of this story - FREE

III. YOUR STOCK TRACKER

Company Name	Ticker	Closing	Change	Volume	52WkHi	52WkLo
AIR & WATER TECH CP CL A	AWT	6.38	_____	21500	8.12	4.00
BANK OF SOUTHINGTON	BSO	13.50	_____	000	19.38	7.00
CABLEVISION SYS CP CL A	CVC	45.62	0.62	27300	69.75	44.00
DEVON ENERGY CP	DVN	24.75	_____	101500	26.12	18.00
EDITEK INC	EDI	1.06	_____	123700	3.94	0.62
FIRST AUSTRALIA PRIME IN	FAX	8.56	0.00	223900	9.88	8.12
GEN MICROWAVE CP	GMW	7.00	_____	000	9.00	5.88
HAWAIIAN AIRLINES INC	HA	4.81	-0.50	76100	11.50	1.62
RF POWER PRODUCTS INC	RFP	6.12	_____	54700	8.38	3.50
VALLEY RESOURCES INC	VR	12.38	_____	200	12.62	10.25
WIRELESS TELECOMM GRP IN	WTT	10.25	-0.38	73900	16.75	6.12
XYTRONYX INC	XYX	2.25	_____	27900	3.69	1.25
ZIEGLER COS INC-WISC	ZCO	18.62	_____	500	20.12	14.75

DOCUMENT ORDERING INSTRUCTIONS // MEMBER SERVICES

D.3.2 Special Edition

Subject: Patrick Schmitt-Heinrich's Special Edition from IBM infoSage
Date: Thu, 27 Jun 1996 02:01:04 -0400
From: sage§chicory.infosage.ibm.com
To: DEIBMC5K§IBMMAIL.COM
Subject: Patrick Schmitt-Heinrich's Special Edition from IBM infoSage

SPORTS SUMMARY - Sports Summary

Reuters Sports News Summary

Sixers Make Iverson No. 1 Pick

The Philadelphia 76ers selected Georgetown's Allen Iverson as the No. 1 pick in the NBA draft. He's the fifth straight underclassmen to be selected first overall in the draft and the first point guard since Magic Johnson was chosen by the Los Angeles Lakers in 1979. The Toronto Raptors held true to their word and made Massachusetts center Marcus Camby the second selection. Camby swept player of the year honors and led UMass to its first-ever Final Four appearance in the NCAA Tournament. The Vancouver Grizzlies selected California forward Shareef Abdur-Rahim with the third pick. The Milwaukee Bucks chose point guard Stephon Marbury of Georgia Tech with the fourth selection. The Minnesota Timberwolves selected guard Ray Allen of Connecticut with the fifth pick.

Seles Ousted, Becker Advances

The rash of upsets at Wimbledon has claimed another victim. Second seed Monica Seles was ousted Wednesday by Katarina Studenikova of Slovakia, 7-5, 5-7, 6-4. The second-round loss was the American's earliest career exit from a Grand Slam tournament. On the men's side, second seed and three-time champion Boris Becker struggled early but settled down to defeat Spain's Tomas Carbonell, 4-6, 6-3, 6-4, 6-2 to advance to the third round. And American MaliVai Washington eliminated ninth seed Thomas Enqvist of Sweden, 6-4, 7-6, 6-3. Enqvist joins top-ten seeds Andre Agassi, Jim Courier and Michael Chang on the sidelines.

Texas Rangers Corral Orioles

The Texas Rangers rallied from a four-run deficit Wednesday for a 6-5 victory over the Baltimore Orioles. Baltimore entered the bottom of the eighth with a 5-3 lead, but three relievers failed to protect it. In other American League action, the New York Yankees edged out Minnesota, 2-1, and extended their lead in the American League East to four games over Baltimore. Toronto sunk Seattle, 6-5; Kansas City downed Milwaukee, 7-3; and Boston took 15 innings to beat Cleveland, 6-4.

Cardinals Bang Out 17 Hits

The St. Louis Cardinals banged out a season-high 17 hits en route to an 11-7 victory Wednesday over the Atlanta Braves. The victory moves St. Louis into first place in the National League Central Division -- a half-game ahead of the Houston Astros. In other National League action, Pittsburgh defeated Montreal, 3-1; the New York Mets downed Colorado, 9-5; Florida beat San Francisco, 3-2; Cincinnati pounded Philadelphia, 4-2; Chicago beat Los Angeles, 6-4; and Houston slid past San Diego, 4-3.

Lasorda Remains Hospitalized

Los Angeles Dodgers manager Tommy Lasorda underwent angioplasty Wednesday to clear a blockage in a coronary artery. The 68-year-old Lasorda drove himself to the hospital Monday

night and was admitted with abdominal pains. Team physician Dr. Mickey Mellman conducted an examination, which determined Lasorda's pain "was felt to be in excess of that which would be explained by his superficial ulcer and additional tests indicated a heart problem," the team announced Wednesday. The angioplasty was performed without complication. Lasorda is listed in stable condition and is resting comfortably, a team spokesman said.

Germany, Czechs in EuroCup Final

Powerhouse Germany will face the underdog Czech Republic in the finals of the EuroCup soccer tournament Sunday in London. Tournament favorite Germany defeated host England, 6-5, Wednesday in penalty kicks. The Czechs won their semifinal match over France by an identical score and also on penalty kicks. The Czechs are the surprise of the tournament. They were 80-1 underdogs to reach the finals. The Czechs may have a glimmer of hope. The Germans will have to play the championship match without star Andreas Moeller, who will be serving a one-game suspension after receiving his second yellow card of the tournament.

Gartner Balks at Trade

Future NHL Hall of Famer Mark Gartner says he may retire rather than report to the Phoenix Coyotes. The fifth-leading scorer in NHL history says the Toronto Maple Leafs broke a verbal agreement when they traded him last week to Phoenix. The 36-year-old right wing says he won't pursue legal action against the Maple Leafs, which he says were looking to trim costs. After 18 years in the league with three different teams, Gartner says he doesn't want to move his family again and may instead hang up his skates.

NHL to Expand?

The National Hockey League, which has added five teams since 1991, will accept applications for new expansion teams, NHL Commissioner Gary Bettman announced Wednesday. Among the factors that will be considered in each application are location, demographics, arena quality and lease terms, media and sponsorship potential, the applicant's financial and management capabilities and other relevant considerations, Bettman said. "Our ultimate goal will be to have new teams that will be both competitively and economically successful, that will add to our fan base and that will enhance the NHL's position in the sports and entertainment marketplace," he added.

Morris Pleads Guilty in Court

Pittsburgh Steelers running back Byron "Bam" Morris pleaded guilty Wednesday to a charge of felony marijuana possession in a Rockwall, Texas, court. As part of a plea bargain, prosecutors will recommend probation instead of jail time for the 24-year-old Morris. The agreement also calls for a separate charge of felony cocaine possession to be dropped. Morris' sentence, which is expected to include community service and a fine, may be decided at a July 11th hearing. He was released

today on $25,000 bail. Morris was arrested March 22 after police found marijuana and cocaine in his Mercedes.

Group Pledges Stadium Money

There's been talk in Milwaukee about the Brewers moving if a new stadium isn't built. To quiet such speculation, the Lynde and Harry Bradley Foundation said Wednesday that it's prepared to invest $20 million in a new stadium. Foundation President Michael Joyce says the group would put up the money to "secure the future of a treasured community asset and preclude the despair and disunity among our citizens that surely would follow the loss of a major league baseball franchise." The Brewers have been struggling for months to borrow $90 million for a proposed $250 million stadium. Miller Brewing Company already has committed to donate $40 million in exchange for naming rights.

Reut02:04 06-27-96

---- End of mail text

D.3.3 Alert Delivery

Subject: VISUAL TELEPHONE GETS MOODY'S LISTING
Date: Wed, 26 Jun 1996 19:02:03 -0400
From: sage§chicory.infosage.ibm.com
To: DEIBMC5K§IBMMAIL.COM
Subject: VISUAL TELEPHONE GETS MOODY'S LISTING

IBM infoSage Alert Delivery

Patrick Schmitt-Heinrich
Member ID: XXXXXXXXX

Company: IBM XXXXXXXXX

Wednesday, 26 June 1996
07:05 PM EDT

Member services information is located at the end of your delivery.

Topic Matched: ATM Network

VISUAL TELEPHONE GETS MOODY'S LISTING

TOTOWA, N.J., June 26 /PRNewswire/ -- Visual Telephone (Nasdaq-Electronic Bulletin Board: VTJB) announced today that it will now be part of Moody's Industrial Manual effective Tuesday, July 2, 1996.

Visual Telephone, who recently purchased the right to hook into the newly-developed multimedia high-bandwidth switched network of IntermediaNet, can now utilize an Asynchronous Transfer Mode (ATM) backbone for transmitting and receiving video, voice, and data signals. By using ATM technology, visual Telephone (VTJB) can transmit these signals at much greater speeds, or bandwidth, then possible with other technology, particularly the most commonly used digital network service, Integrated Services Digital Network ("ISDN").

President, Joel Beagelman, stated that, "This is only the beginning" for the company is about to push off its video conferencing centers through both franchising and corporate ownership.

-O- 6/26/96 /CONTACT: Wall Street Associates, 516-889-0163/ (VTJB)
CO: Visual Telephone ST: New Jersey IN: TLS SU: RTG

Appendix E. IBM Global Network Phone List

The following are the 600 access numbers to the IBM Global Network that are currently around the world.

- Argentina Buenos Aires 319-7202
- Australia Adelaide 08-357-8794
- Australia Ballarat 053-302-915
- Australia Brisbane 07-3832-9188
- Australia Canberra 06-273-5269
- Australia Darwin 08-8981-3933
- Australia Hobart 002-248-391
- Australia Melbourne (V.34) 03-9690-3300
- Australia Newcastle 049-262 287
- Australia Perth 09-321-7199
- Australia Sydney (V.34) 02-899-3399
- Australia Wollongong 042 296 955
- Austria Bregenz 05574-43875
- Austria Eisenstadt 02682-72250
- Austria Graz 0316-915096
- Austria Innsbruck 0512-579549
- Austria Klagenfurt 0463-511924
- Austria Linz 0732-783615
- Austria Salzburg 0662-827692
- Austria St. Poelten (V.34) 02742-71720
- Austria Vienna-DON (V.34) 0222-2162610
- Austria Vienna-LAS (V.34) 0222-2144020
- Belgium Antwerpen 03-2486565
- Belgium Brussels (V.34) 02-7209291
- Belgium Gent 09-2210674
- Belgium Liège 041-672686
- Brazil Fortaleza (085)255-0505
- Brazil Rio de Janeiro (021)516-2020
- Brazil Salvador (071)353-4466
- Brazil Sao Paulo (011)870-5757
- Brazil Sao Paulo (011)885-7799
- Bulgaria Sofia 02 71463094
- Canada AB Calgary (V.34) (403) 290-5651
- Canada PE Charlottetown (V.34) (902) 629-4659

© Copyright IBM Corp. 1996

- Canada AB Edmonton (V.34) (403) 917-4451
- Canada NS Halifax (V.34) (902) 493-1321
- Canada ON Hamilton (V.34) (905) 540-2551
- Canada ON London (V.34) (519) 640-8401
- Canada PQ Montreal (V.34) (514) 846-7171
- Canada ON Ottawa (V.34) (613) 788-0706
- Canada PQ Quebec City (V.34) (418) 525-3101
- Canada SK Regina (V.34) (306) 566-7501
- Canada NB Saint John (V.34) (506) 658-3581
- Canada NF St. John's (V.34) (709) 570-8801
- Canada ON Toronto (V.34) (416) 758-5871
- Canada BC Vancouver (V.34) (604) 602-2401
- Canada BC Victoria (V.34) (604) 995-3751
- Canada ON Waterloo (V.34) (519) 885-8001
- Canada MB Winnipeg (V.34) (204) 934-6301
- Canada ON Windsor (V.34) (519) 972-4541
- Canada fee 800 (V.34) 1 (800) 250-6333
- Chile Santiago (2) 2351132
- Chile (2) 2363750
- Chile (2) 2006555
- Colombia Bogota 571-256-9311
- Cyprus Nicosia 1601
- Czech Republic Brno 05 43215495
- Czech Republic Ceske Budejovice 038 28643
- Czech Republic Hradec Kralove 049 617205
- Czech Republic Karlovy Vary 017 3221512
- Czech Republic Olomouc 068 91232
- Czech Republic Ostrava 069 51556
- Czech Republic Plzen 019 7235659
- Czech Republic Prague 02 67106408
- Czech Republic Usti Nad Labem 047 5200935
- Czech Republic Zlin 067 31512
- Denmark Aarhus 8739 6060
- Denmark Copenhagen (V.34) 4593 4290
- Ecuador Quito 528-033
- Finland Helsinki (V.34) 90-4587100
- France Bordeaux 56 69 97 25
- France Lille 20 24 04 48

- France Lyon 78 33 63 73
- France Marseille 91 77 03 56
- France Montpellier 67 22 34 25
- France Nantes 40 47 19 63
- France Nice (La Gaude) 92 11 02 02
- France Orleans 38 55 28 00
- France Paris-East (V.34) 161-43052770
- France Paris-West (V.34) 161-47784307
- France Strasbourg 88 25 66 46
- France Toulouse 61 20 73 57
- Germany Augsburg 0821-3493800
- Germany Bayreuth 0921-560498
- Germany Berlin 030-7231021
- Germany Berlin (V.34) 030-7231331
- Germany Bremen 0421-2439958
- Germany Chemnitz 0371-306567
- Germany Dresden 0351-4903571
- Germany Duesseldorf 0211-432155
- Germany Ehningen (V.34) 07034-93600
- Germany Erfurt 0361-6442450
- Germany Essen 0201-7109100
- Germany Frankfurt 069-6668542
- Germany Frankfurt (V.34) 069-6613011
- Germany Freiburg 0761-2020932
- Germany Hamburg (V.34) 040-6303655
- Germany Hannover 0511-9524744
- Germany Karlsruhe 0721-892180
- Germany Kassel 0561-780822
- Germany Kiel 0431-641925
- Germany Koblenz 0261-16204
- Germany Koeln 0221-3405026
- Germany Leipzig 0341-9608340
- Germany Magdeburg 0391-5410800
- Germany Mainz 06131-834630
- Germany Mannheim 0621-401026
- Germany Munich (V.34) 089-342418
- Germany Muenster 0251-2305160
- Germany Nuernberg 0911-813043

- Germany Regensburg 0941-792344
- Germany Rostock 0381-4004800
- Germany Saarbruecken 0681-31362
- Germany Stuttgart 0711-7800264
- Germany Ulm 0731-6020700
- Germany Wuerzburg 0931-781936
- Greece Athens 01-6801330
- Greece Thessaloniki 031-244540
- Hong Kong (V.34) 3004-9009
- Hong Kong (backup) 3004-9600
- Hungary Budapest 1 185 2627
- Indonesia Bandung 022-7277070
- Indonesia Jakarta (V.34) 021-5270870
- Indonesia Jakarta (V.34) 021-5209580
- Indonesia Jakarta 021-3507070
- Indonesia Medan 061-547070
- Ireland Dublin 01-6607100
- Israel Haifa 04-8550696
- Israel Tel Aviv 03-695-2777
- Italy Bari 080-54.14.108
- Italy Bologna 051-64.14.300
- Italy Firenze 055-32.00.110
- Italy Genova 010-57.00.420
- Italy Milano (V.34) 02 70300693
- Italy Napoli 081-22.20.303
- Italy Padova 049-666311
- Italy Palermo 091-61.19.360
- Italy Roma 06-59648366
- Italy Torino 011-7777870
- Italy Verona 045-82.67.120
- Japan nationwide (V.34) 0088-36-1111
- Japan Chiba 043-285-5681
- Japan Fukuoka 092-621-7942
- Japan Hiroshima 082-225-3555
- Japan Kawasaki 044-245-7604
- Japan Kobe 078-612-3792
- Japan Nagoya 052-581-9571
- Japan Niigata 025-245-9199

- Japan Ohmiya 048-647-7007
- Japan Osaka 06-376-3027
- Japan Sapporo 011-752-9732
- Japan Sendai 022-262-7421
- Japan Tokyo 03-3505-5885
- Japan Tokyo (V.34) 03-3505-4891
- Japan Yokohama 045-451-2431
- Luxembourg 366957
- Malaysia Kuala Lumpur 7162516
- Malaysia Kuala Lumpur 7162554
- Mexico Guadalajara 52(3)689 0807
- Mexico Mexico City 52(5)327 5850
- Mexico Monterrey 52(8)319 0633
- Mexico Tijuana 52(66)34 1060
- N.A. Curacao 369-811
- N.A. St. Marten 0255
- Netherlands Amsterdam (V.34) 020-6151500
- Netherlands Arnhem 026-3888062
- Netherlands Eindhoven 040-2465808
- Netherlands Groningen 050-5260022
- Netherlands Utrecht 030-2804844
- Netherlands Zoetermeer (V.34) 079-3212244
- Netherlands Zwolle 038-4235500
- New Zealand Auckland 09-356-3984
- New Zealand Christchurch 03-372-8954
- New Zealand Wellington 04-576-5998
- Norway Oslo, for Oslo and Akershus (V.34) 66809022
- Norway Oslo, for rest of Norway (V.34) 81003555
- Peru Lima (01) 349-0165
- Philippines Manila (02) 8126060
- Portugal Lisbon 01 7915145
- Portugal Porto 02 2071145
- Russian Federation Moscow 095-258-6420
- Slovakia Bratislava 07 787931
- Slovenia Ljubljana 061 1264777
- Slovenia Murska Sobota 069 27-075
- South Africa Cape Town 021-4013380
- South Africa Durban 031-2682380

- South Africa Johannesburg 011-7001188
- South Africa Johannesburg 011-7001766
- Spain Alicante 96-5116539
- Spain Barcelona 93-3220000
- Spain Barcelona 93-3220505
- Spain Bilbao 94-4167000
- Spain Bilbao 94-4157922
- Spain La Coruna 981-226388
- Spain Las Palmas 928-383688
- Spain Madrid (V.34) 91-6567258
- Spain Madrid (V.34) 91-6567105
- Spain Madrid (V.34) 91-4137314
- Spain Madrid (V.34) 91-4130011
- Spain Malaga 95-2228800
- Spain Murcia 968-280228
- Spain Oviedo 98-5275755
- Spain Palma de Mallorca 971-755195
- Spain Pamplona 948-236689
- Spain San Sebastian 943-217577
- Spain Sevilla 95-4280710
- Spain Sevilla 95-4282960
- Spain Sta.Cruz de Tenerife 922-243288
- Spain Valencia 96-3930190
- Spain Valencia 96-3933355
- Spain Vigo 986-231211
- Spain Zaragoza 976-212018
- Sweden Gothenburg 031 80 21 13
- Sweden Malmoe 040 12 31 15
- Sweden Stockholm (V.34) 08-6320040
- Switzerland Basel 061 2740100
- Switzerland Basel (V34) 061 2741101
- Switzerland Bern (V34) 031 3827070
- Switzerland Chur (V34) 081 2528334
- Switzerland Geneva (V34) 022 7336633
- Switzerland Lausanne (V34) 021 3129010
- Switzerland Lugano (V34) 091 9233442
- Switzerland Luzern (V34) 041 2100133
- Switzerland Olten (V34) 062 2960130

- Switzerland St.Gallen (V34) 071 2232222
- Switzerland Zurich 01 4330320
- Switzerland Zurich (V34) 01 8031800
- Taiwan 7786565
- Thailand 236-9652
- Turkiye Ankara 0468-3535
- Turkiye Ankara 0468-3039
- Turkiye Ankara 0468-2531
- Turkiye Ankara 0468-3492
- Turkiye Istanbul 0212-2823944
- UK Bristol 0117-9292037
- UK Edinburgh 0131-5570465
- UK Glasgow 0141-226-4659
- UK Leeds 01132-433878
- UK London (Greenford) 0181-575-7633
- UK London SBK (V.34) 0171-6203415
- UK Manchester 0161-9621452
- UK Nottingham 0115-9419214
- UK Portsmouth (V.34) 01705-325027
- UK Warwick (V.34) 01926-493401
- US AK Anchorage (V.34) 1 (907) 343-3501
- US AL Birmingham (V.34) 1 (205) 510-2001
- US AL Dothan (V.34) 1 (334) 615-9001
- US AL Florence (V.34) 1 (205) 760-3361
- US AL Huntsville (V.34) 1 (205) 890-1901
- US AL Mobile (V.34) 1 (334) 602-6501
- US AL Montgomery (V.34) 1 (334) 409-4601
- US AL Tuscaloosa (V.34) 1 (205) 391-0693
- US AR Fayetteville (V.34) 1 (501) 587-6601
- US AR Fort Smith (V.34) 1 (501) 452-0290
- US AR Little Rock (V.34) 1 (501) 791-8701
- US AR Pine Bluff (V.34) 1 (501) 541-5901
- US AZ Phoenix (V.34) 1 (602) 395-5301
- US AZ Tucson (V.34) 1 (602) 512-1201
- US CA Bakersfield (V.34) 1 (805) 396-3901
- US CA Capistrano (V.34) 1 (714) 460-7101
- US CA Chico (V.34) 1 (916) 891-7701
- US CA Concord (V.34) 1 (510) 687-0138

- US CA Costa Mesa (V.34) 1 (714) 825-1301
- US CA Escondido (V.34) 1 (619) 735-5301
- US CA Fresno (V.34) 1 (209) 248-4101
- US CA Lancaster (V.34) 1 (805) 726-3441
- US CA Los Angeles (V.34) 1 (213) 893-9501
- US CA Marysville (V.34) 1 (916) 749-7801
- US CA Merced (V.34) 1 (209) 384-5301
- US CA Monterey (V.34) 1 (408) 645-7401
- US CA Morgan Hill (V.34) 1 (408) 776-7201
- US CA Napa (V.34) 1 (707) 254-1005
- US CA Norwalk (V.34) 1 (310) 497-4401
- US CA Ontario (V.34) 1 (909) 930-2001
- US CA Palo Alto (V.34) 1 (415) 846-5901
- US CA Redding (V.34) 1 (916) 242-4301
- US CA Riverside (V.34) 1 (909) 341-0998
- US CA Sacramento (V.34) 1 (916) 863-9501
- US CA Salinas (V.34) 1 (408) 442-6701
- US CA San Bernardino (V.34) 1 (909) 888-0190
- US CA San Diego (V.34) 1 (619) 657-5501
- US CA San Francisco (V.34) 1 (415) 827-2201
- US CA San Jose (V.34) 1 (408) 289-5701
- US CA San Ramon (V.34) 1 (510) 867-0544
- US CA Santa Barbara (V.34) 1 (805) 737-3401
- US CA Santa Cruz (V.34) 1 (408) 477-4601
- US CA Stockton (V.34) 1 (209) 475-4401
- US CA Ventura (V.34) 1 (805) 383-4301
- US CA Victorville (V.34) 1 (619) 381-8301
- US CA Visalia (V.34) 1 (209) 741-2301
- US CA Woodland Hills (V.34) 1 (818) 595-0018
- US CO Boulder (V.34) 1 (303) 605-2101
- US CO Colorado Sprgs (V.34) 1 (719) 527-3941
- US CO Grand Junction (V.34) 1 (970) 256-8201
- US CO Pueblo (V.34) 1 (719) 585-1601
- US CT Danbury (V.34) 1 (203) 207-3001
- US CT Fairfield (V.34) 1 (203) 319-2401
- US CT Hamden (V.34) 1 (203) 781-1601
- US CT Hartford (V.34) 1 (203) 550-7201
- US CT Milford (V.34) 1 (203) 876-1285

- US CT New Haven (V.34) 1 (203) 781-1601
- US CT New London (V.34) 1 (203) 405-2001
- US CT Norwalk (V.34) 1 (203) 845-0623
- US CT Stamford (V.34) 1 (203) 348-0021
- US CT Waterbury (V.34) 1 (203) 262-7301
- US DC Washington (V.34) 1 (301) 754-3901
- US DE Wilmington (V.34) 1 (302) 425-0116
- US FL Boca Raton (V.34) 1 (407) 447-2001
- US FL Daytona Beach (V.34) 1 (904) 947-5401
- US FL Fort Lauderdale (V.34) 1 (954) 771-1343
- US FL Fort Myers (V.34) 1 (813) 277-3761
- US FL Gainesville (V.34) 1 (352) 333-7001
- US FL Jacksonville (V.34) 1 (904) 419-2501
- US FL Lakeland (V.34) 1 (941) 499-1601
- US FL Miami (V.34) 1 (305) 460-9501
- US FL Ocala (V.34) 1 (352) 694-9001
- US FL Orlando (V.34) 1 (407) 673-3901
- US FL Panama City (V.34) 1 (904) 913-7301
- US FL Pensacola (V.34) 1 (904) 969-3001
- US FL Sarasota (V.34) 1 (941) 331-4101
- US FL St. Petersburg (V.34) 1 (813) 524-7101
- US FL Tallahassee (V.34) 1 (904) 216-0901
- US FL Tampa (V.34) 1 (813) 554-1101
- US FL Vero Beach (V.34) 1 (407) 564-6141
- US FL West Palm Beach (V.34) 1 (407) 615-5701
- US GA Albany (V.34) 1 (912) 430-2601
- US GA Athens (V.34) 1 (706) 613-7371
- US GA Atlanta (V.34) 1 (770) 270-6901
- US GA Augusta (V.34) 1 (706) 739-1001
- US GA Columbus (V.34) 1 (706) 562-1551
- US GA Macon (V.34) 1 (912) 757-5801
- US GA Savannah (V.34) 1 (912) 692-4421
- US HI Honolulu (V.34) 1 (808) 979-5101
- US IA Cedar Falls (V.34) 1 (319) 236-6901
- US IA Cedar Rapids (V.34) 1 (319) 395-6601
- US IA Des Moines (V.34) 1 (515) 267-6801
- US IA Dubuque (V.34) 1 (319) 557-5801
- US IA Sioux City (V.34) 1 (712) 274-6201

- US ID Boise (V.34) 1 (208) 333-2501
- US ID Coeur D'alene (V.34) 1 (208) 667-0955
- US ID Idaho Falls (V.34) 1 (208) 535-0101
- US IL Bloomington (V.34) 1 (309) 664-2401
- US IL Champaign (V.34) 1 (217) 351-1301
- US IL Chicago (V.34) 1 (312) 464-6251
- US IL Chicago Ridge (V.34) 1 (708) 229-8001
- US IL Decatur (V.34) 1 (217) 421-1501
- US IL Elmhurst (V.34) 1 (708) 613-1201
- US IL Moline (V.34) 1 (319) 388-5401
- US IL Peoria (V.34) 1 (309) 694-8201
- US IL Rockford (V.34) 1 (815) 332-6701
- US IL Schaumburg (V.34) 1 (708) 237-1101
- US IL Springfield (V.34) 1 (217) 793-7181
- US IN Anderson (V.34) 1 (317) 683-4301
- US IN Bloomington (V.34) 1 (812) 349-4101
- US IN Evansville (V.34) 1 (812) 469-7001
- US IN Fort Wayne (V.34) 1 (219) 470-9001
- US IN Indianapolis (V.34) 1 (317) 655-3001
- US IN Kokomo (V.34) 1 (317) 864-6201
- US IN Lafayette (V.34) 1 (317) 429-5436
- US IN Marion (V.34) 1 (317) 677-5001
- US IN Merrillville (V.34) 1 (219) 681-6401
- US IN Muncie (V.34) 1 (317) 747-1115
- US IN South Bend (V.34) 1 (219) 271-2601
- US IN Terre Haute (V.34) 1 (812) 231-8701
- US KS Lawrence (V.34) 1 (913) 838-0201
- US KS Manhattan (V.34) 1 (913) 565-3001
- US KS Salina (V.34) 1 (913) 452-3101
- US KS Topeka (V.34) 1 (913) 228-8301
- US KS Wichita (V.34) 1 (316) 337-9501
- US KY Lexington (V.34) 1 (606) 245-7201
- US KY Louisville (V.34) 1 (502) 499-4018
- US KY Owensboro (V.34) 1 (502) 688-9401
- US KY Paducah (V.34) 1 (502) 575-9603
- US LA Alexandria (V.34) 1 (318) 483-3901
- US LA Baton Rouge (V.34) 1 (504) 930-3001
- US LA Lafayette (V.34) 1 (318) 983-7201

- US LA Lake Charles (V.34) 1 (318) 437-1801
- US LA Monroe (V.34) 1 (318) 329-0101
- US LA New Orleans (V.34) 1 (504) 456-3641
- US LA Shreveport (V.34) 1 (318) 226-7001
- US LA Slidell (V.34) 1 (504) 639-3301
- US MA Boston (V.34) 1 (617) 927-5101
- US MA Lawrence (V.34) 1 (508) 837-6301
- US MA Lexington (V.34) 1 (617) 276-4101
- US MA New Bedford (V.34) 1 (508) 646-4201
- US MA Springfield (V.34) 1 (413) 543-7601
- US MA Worcester (V.34) 1 (508) 890-2601
- US MD Annapolis (V.34) 1 (410) 216-7801
- US MD Baltimore (V.34) 1 (410) 771-8981
- US MD Cumberland (V.34) 1 (301) 729-5901
- US MD Hagerstown (V.34) 1 (301) 665-1501
- US MD Salisbury (V.34) 1 (410) 334-3001
- US ME Augusta (V.34) 1 (207) 626-5101
- US ME Bangor (V.34) 1 (207) 990-0614
- US ME Lewiston (V.34) 1 (207) 753-2501
- US ME Portland (V.34) 1 (207) 842-5201
- US MI Ann Arbor (V.34) 1 (313) 913-8112
- US MI Battle Creek (V.34) 1 (616) 963-9949
- US MI Detroit Downtown (V.34) 1 (313) 202-1101
- US MI Detroit Southfield (V.34) 1 (810) 204-1301
- US MI Flint (V.34) 1 (810) 733-9441
- US MI Grand Rapids (V.34) 1 (616) 975-1601
- US MI Jackson (V.34) 1 (517) 796-6001
- US MI Kalamazoo (V.34) 1 (616) 341-4749
- US MI Lansing (V.34) 1 (517) 333-9743
- US MI Midland (V.34) 1 (517) 832-0603
- US MI Muskegon (V.34) 1 (616) 728-8506
- US MI Novi (V.34) 1 (810) 347-7401
- US MI Saginaw (V.34) 1 (517) 249-1901
- US MI St. Joseph (V.34) 1 (616) 428-0702
- US MI Traverse City (V.34) 1 (616) 922-0126
- US MN Duluth (V.34) 1 (218) 725-0001
- US MN Mankato (V.34) 1 (507) 386-4601
- US MN Minneapolis (V.34) 1 (612) 943-5801

- US MN Rochester (V.34) 1 (507) 287-9681
- US MN St. Cloud (V.34) 1 (612) 202-2201
- US MO Cape Girardeau (V.34) 1 (314) 986-6601
- US MO Columbia (V.34) 1 (573) 499-9581
- US MO Jefferson City (V.34) 1 (573) 556-6001
- US MO Joplin (V.34) 1 (417) 659-5401
- US MO Kansas City (V.34) 1 (816) 795-3101
- US MO Rolla (V.34) 1 (314) 364-6372
- US MO Springfield (V.34) 1 (417) 891-1901
- US MO St. Joseph (V.34) 1 (816) 236-1101
- US MO St. Louis (V.34) 1 (314) 551-1501
- US MS Biloxi (V.34) 1 (601) 385-5101
- US MS Gulfport (V.34) 1 (601) 863-9728
- US MS Hattiesburg (V.34) 1 (601) 543-1101
- US MS Jackson (V.34) 1 (601) 346-1001
- US MS Meridian (V.34) 1 (601) 481-2001
- US MS Tupelo (V.34) 1 (601) 840-6911
- US MT Billings (V.34) 1 (406) 238-4741
- US MT Bozeman (V.34) 1 (406) 582-7301
- US MT Butte (V.34) 1 (406) 494-8611
- US MT Great Falls (V.34) 1 (406) 771-4181
- US MT Helena (V.34) 1 (406) 443-8101
- US MT Missoula (V.34) 1 (406) 542-6301
- US NC Asheville (V.34) 1 (704) 299-5201
- US NC Charlotte (V.34) 1 (704) 510-2001
- US NC Durham (V.34) 1 (919) 558-5901
- US NC Fayetteville (V.34) 1 (910) 860-4001
- US NC Greensboro (V.34) 1 (910) 605-1541
- US NC Greenville (V.34) 1 (919) 353-1801
- US NC High Point (V.34) 1 (910) 881-3801
- US NC Raleigh (V.34) 1 (919) 878-4801
- US NC Rocky Mount (V.34) 1 (919) 407-1461
- US NC Wilmington (V.34) 1 (910) 792-5341
- US NC Winston-Salem (V.34) 1 (910) 733-5971
- US ND Bismarck (V.34) 1 (701) 250-0001
- US ND Fargo (V.34) 1 (701) 281-6401
- US ND Grand Forks (V.34) 1 (701) 795-4701
- US ND Minot (V.34) 1 (701) 858-1901

- US NE Grand Island (V.34) 1 (308) 385-1501
- US NE Lincoln (V.34) 1 (402) 467-9501
- US NE Omaha (V.34) 1 (402) 392-5101
- US NH Manchester (V.34) 1 (603) 634-0801
- US NJ Cherry Hill (V.34) 1 (609) 488-1919
- US NJ Mays Landing (V.34) 1 (609) 569-7201
- US NJ Paramus (V.34) 1 (201) 986-2741
- US NJ Princeton (V.34) 1 (609) 514-7501
- US NJ Toms River (V.34) 1 (908) 473-1085
- US NJ Trenton (V.34) 1 (609) 278-7701
- US NJ West Orange (V.34) 1 (201) 325-4401
- US NM Albuquerque (V.34) 1 (505) 837-8101
- US NM Las Cruces (V.34) 1 (505) 523-5621
- US NM Santa Fe (V.34) 1 (505) 438-3806
- US NV Las Vegas (V.34) 1 (702) 693-5601
- US NV Reno (V.34) 1 (702) 785-9001
- US NY Albany (V.34) 1 (518) 454-3301
- US NY Buffalo (V.34) 1 (716) 568-8301
- US NY Corning (V.34) 1 (607) 796-2135
- US NY Endicott (V.34) 1 (607) 766-1001
- US NY Ithaca (V.34) 1 (607) 275-1001
- US NY Jamestown (V.34) 1 (716) 665-1401
- US NY Jericho (V.34) 1 (516) 733-3561
- US NY Kingston (V.34) 1 (914) 334-2601
- US NY Lockport (V.34) 1 (716) 433-0071
- US NY New York City (V.34) 1 (212) 605-5101
- US NY Poughkeepsie (V.34) 1 (914) 431-1281
- US NY Rochester (V.34) 1 (716) 246-4001
- US NY Syracuse (V.34) 1 (315) 448-1101
- US NY Utica (V.34) 1 (315) 793-2921
- US NY White Plains (V.34) 1 (914) 683-6001
- US OH Akron (V.34) 1 (216) 342-1301
- US OH Canton (V.34) 1 (216) 492-3391
- US OH Cincinnati (V.34) 1 (513) 741-6581
- US OH Cleveland (V.34) 1 (216) 843-4481
- US OH Columbus (V.34) 1 (614) 272-4201
- US OH Dayton (V.34) 1 (513) 438-7101
- US OH Lima (V.34) 1 (419) 221-5301

- US OH Mansfield (V.34) 1 (419) 521-0001
- US OH Newark (V.34) 1 (614) 323-5001
- US OH Smithfield (V.34) 1 (614) 284-5501
- US OH Springfield (V.34) 1 (513) 324-9641
- US OH Toledo (V.34) 1 (419) 244-3085
- US OH Youngstown (V.34) 1 (216) 629-7881
- US OK Enid (V.34) 1 (405) 548-0401
- US OK Oklahoma City (V.34) 1 (405) 280-2501
- US OK Tulsa (V.34) 1 (918) 488-1201
- US OR Bend (V.34) 1 (541) 383-8910
- US OR Corvallis (V.34) 1 (541) 757-1748
- US OR Eugene (V.34) 1 (541) 485-0102
- US OR Medford (V.34) 1 (541) 857-4501
- US OR Portland (V.34) 1 (503) 223-0904
- US OR Salem (V.34) 1 (503) 373-9619
- US PA Bethleham (V.34) 1 (610) 807-4321
- US PA Erie (V.34) 1 (814) 866-4401
- US PA Harrisburg (V.34) 1 (717) 671-2601
- US PA Lancaster (V.34) 1 (717) 581-5921
- US PA Philadelphia (V.34) 1 (215) 851-8301
- US PA Pittsburgh (V.34) 1 (412) 237-2301
- US PA Reading (V.34) 1 (610) 208-4561
- US PA Scranton (V.34) 1 (717) 340-6781
- US PA State College (V.34) 1 (814) 238-0380
- US PA Wilkes-Barre (V.34) 1 (717) 831-0701
- US PA Williamsport (V.34) 1 (717) 327-7481
- US PA York (V.34) 1 (717) 771-1001
- US PR Santurce (V.34) 1 (809) 289-0801
- US RI Providence (V.34) 1 (401) 827-5401
- US SC Charleston (V.34) 1 (803) 820-4601
- US SC Columbia (V.34) 1 (803) 865-6101
- US SC Florence (V.34) 1 (803) 317-0001
- US SC Greenville (V.34) 1 (803) 234-2001
- US SC Myrtle Beach (V.34) 1 (803) 444-1301
- US SC Spartanburg (V.34) 1 (803) 515-5761
- US SD Sioux Falls (V.34) 1 (605) 373-3201
- US TN Chattanooga (V.34) 1 (423) 954-3901
- US TN Clarksville (V.34) 1 (615) 905-5701

- US TN Cleveland (V.34) 1 (423) 559-0223
- US TN Jackson (V.34) 1 (901) 423-7601
- US TN Johnson City (V.34) 1 (423) 975-7701
- US TN Kingsport (V.34) 1 (423) 392-9812
- US TN Knoxville (V.34) 1 (423) 595-4601
- US TN Memphis (V.34) 1 (901) 762-9801
- US TN Nashville (V.34) 1 (615) 731-9931
- US TX Abilene (V.34) 1 (915) 690-4301
- US TX Amarillo (V.34) 1 (806) 354-3801
- US TX Austin (V.34) 1 (512) 302-6501
- US TX Beaumont (V.34) 1 (409) 723-1852
- US TX Corpus Christi (V.34) 1 (512) 994-7101
- US TX Dallas (V.34) 1 (214) 780-2201
- US TX El Paso (V.34) 1 (915) 783-3101
- US TX Fort Worth (V.34) 1 (817) 570-4301
- US TX Harlingen (V.34) 1 (210) 430-1101
- US TX Houston (V.34) 1 (713) 897-6201
- US TX Laredo (V.34) 1 (800) 650-8839
- US TX Longview (V.34) 1 (903) 232-2901
- US TX Lubbock (V.34) 1 (806) 788-2601
- US TX McAllen (V.34) 1 (210) 631-2319
- US TX Midland (V.34) 1 (915) 688-0801
- US TX Odessa (V.34) 1 (915) 368-3001
- US TX San Angelo (V.34) 1 (915) 947-4101
- US TX San Antonio (V.34) 1 (210) 242-7301
- US TX Texarkana (V.34) 1 (903) 794-8241
- US TX Tyler (V.34) 1 (903) 579-7901
- US TX Victoria (V.34) 1 (512) 582-5301
- US TX Waco (V.34) 1 (817) 751-3901
- US TX Wichita Falls (V.34) 1 (817) 696-6801
- US UT Provo (V.34) 1 (801) 344-5001
- US UT Salt Lake City (V.34) 1 (801) 321-6201
- US VA Charlottesville (V.34) 1 (804) 974-5701
- US VA Fredricksburg (V.34) 1 (540) 374-0501
- US VA Lynchburg (V.34) 1 (804) 237-8601
- US VA Manassas (V.34) 1 (703) 361-0018
- US VA Norfolk (V.34) 1 (804) 473-5401
- US VA Petersburg (V.34) 1 (804) 863-4101

- US VA Richmond (V.34) 1 (804) 674-1501
- US VA Roanoke (V.34) 1 (540) 776-6101
- US VA Williamsburg (V.34) 1 (804) 259-5701
- US VT Burlington (V.34) 1 (802) 651-5401
- US WA Bellingham (V.34) 1 (360) 715-7701
- US WA Kennewick (V.34) 1 (509) 783-1895
- US WA Port Angeles (V.34) 1 (360) 417-1501
- US WA Seattle (V.34) 1 (206) 344-3401
- US WA Spokane (V.34) 1 (509) 484-6101
- US WA Tacoma (V.34) 1 (206) 620-1601
- US WA Yakima (V.34) 1 (509) 966-9799
- US WI Appleton (V.34) 1 (414) 830-4501
- US WI Eau Claire (V.34) 1 (715) 831-4001
- US WI Green Bay (V.34) 1 (414) 430-4201
- US WI Janesville (V.34) 1 (608) 368-2101
- US WI Lacrosse (V.34) 1 (608) 796-2401
- US WI Madison (V.34) 1 (608) 243-1601
- US WI Milwaukee (V.34) 1 (414) 860-4201
- US WI Racine (V.34) 1 (414) 554-5480
- US WI Sheboygan (V.34) 1 (414) 451-5901
- US WI Wausau (V.34) 1 (715) 843-2501
- US WV Bridgeport (V.34) 1 (304) 848-0001
- US WV Charleston (V.34) 1 (304) 340-1901
- US WV Clarksburg (V.34) 1 (304) 626-1001
- US WV Huntington (V.34) 1 (304) 697-2364
- US WV Parkersburg (V.34) 1 (304) 420-7801
- US WY Casper (V.34) 1 (307) 261-4101
- US WY Cheyenne (V.34) 1 (307) 637-9101
- US Fee 800 1 (800) 933-3997
- US Fee 800 (V.34) 1 (800) 590-4857
- Venezuela 9088960

For updated information, refer to:

- http://www.ibm.net/phoneint.html

Appendix F. IBM Global Network Registration Phone List

The following is a list of local dial numbers for online registration for the IBM Global Network in each country.

- Argentina Registration 319-7201
- Australia Registration 1800-811-094
- Austria Registration 0660-6832
- Belgium Registration 0800-1-1997
- Brazil Fortaleza Registration (085)255-0505
- Brazil Rio de Janeiro Registration (021)516-2020
- Brazil Sao Paulo Registration (011)885-7799
- Bulgaria Registration 0031 297 532050
- Canada Registration 1-800-463-8331
- Colombia (in Bogota) Registration 571-6167555
- Colombia (outside Bogota) Registration 9800-17555
- Curacao N.A. Registration 368-039
- Cyprus Registration 080-91027
- Czech Republic Registration 0031 297 532050
- Denmark Registration 8001-8278
- Ecuador Registration 565-090
- Finland Registration 0800-114465
- France Registration 0590-8561
- Germany Registration 0130-821202
- Greece Registration 00800-4412-2357
- Hong Kong Registration 2515-2434
- Hungary Registration 0031 297 532050
- Ireland Registration 1-800-709-905
- Indonesia Registration 021-5223431
- Italy Registration 1678-72031
- Israel Registration 177-440-6299
- Japan Registration 0120-120-208
- Luxembourg Registration 0800-2943
- Malaysia Registration 7054500
- Mexico City Registration 52(5) 627 2444
- Netherlands Registration 060-228488
- New Zealand Registration 0800-105765
- Norway Registration 800-11783

© Copyright IBM Corp. 1996

- Peru Manual Registration use extension 1760 (511)-349-0050
- Philippines Registration (632) 8436917
- Philippines Registration (632) 8436918
- Slovak Republic Registration 0031 297 532050
- South Africa Registration 0800-998128
- Spain Registration 900-994443
- Sweden Registration 020-795181
- Switzerland Registration 155-9222
- Taiwan Registration 7786565
- Thailand Registration 001-61-2-894-5166
- Turkey Registration 00800-44914835
- United Kingdom Registration 0800-614012
- United States Registration 1-800-933-3997

For updated information, refer to:

- http://www.ibm.net/phoneint.html

Appendix G. IBM Global Network Help Desk Phone List

The following is a list of local dial numbers for the Internet help desk for the IBM Global Network in each country.

- Argentina 313-0014
- Australia 131-426
- Austria 0660-5702
- Belgium (Dutch) 0800-13270
- Belgium 0800-16521
- Brazil 011-885-0080
- Bulgaria 0031-79-3224516
- Canada (English) 1-800-821-4612
- Chile 800-203037
- Colombia 571-623-2300
- Curacao N.A. 370-360
- Cyprus 080-92205
- Czech Republic 0031-79-3224517
- Denmark 8001-8299
- Finland 0800-1-13151
- France 05-906088
- Germany 0130-821141
- Greece 30-1-3281421
- Hong Kong (852)2515-4511
- Hungary 00800-11516
- Ireland 1-800-553175
- Indonesia 62-21-5238491
- Italy 1678-76007
- Japan 0422-42-8625
- Luxembourg (French) 0800-2921
- Luxembourg (German) 0800-2922
- Malaysia 03-719-2200
- Mexico 91-800-50-567
- Mexico City 52-5-327-5737
- Netherlands 060-222308
- New Zealand 0800-801-800
- Norway 800-11341
- Peru (511)-349-0050

© Copyright IBM Corp. 1996

- Philippines (632) 8192277
- Slovak Republic 0031-79-3224521
- South Africa 011-7001370
- South Africa 0800-117888
- South Africa 0800-110756
- Spain 900-993150
- Sweden 020-795701
- Switzerland (French) 155-9169
- Switzerland (German) 155-9170
- Switzerland (Italian) 155-9173
- Taiwan 7767700
- Thailand 273-4347
- Turkey 90-212-2800900x3305
- United Kingdom 0800-963949
- United States 1-800-821-4612
- Venezuela 800-DEIBM (800-33426)

For updated information, refer to:

- http://www.ibm.net/helpdesk.html?155,60

Appendix H. Special Notices

This publication is intended to help Customers, IBM technical professionals, service specialists, marketing specialists and marketing representatives to define and design a complete solution for the Internet environment.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM′s product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM′s intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY 10594 USA.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM (″vendor″) products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer′s ability to evaluate and integrate them into the customer′s operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

ADSTAR	AIX
AIX/6000	AnyNet
Application System/400	APPN
AS/400	AT
BookManager	BookMaster
CICS	CICS/6000
Cryptolope	DatagLANce
DB2	DB2/2
DB2/6000	DRDA

© Copyright IBM Corp. 1996

EtherStreamer	GDDM
Global Network	Hyperwise
IBM	IMS
ISSC	LAN Distance
LANStreamer	MQ
MQSeries	MVS/ESA
NetDoor	NetFinity
NetView	Nways
Operating System/2	Operating System/400
OS/2	OS/390
OS/400	Personal System/2
Power Series	PowerPC
POWERparallel	PS/1
PS/2	RACF
RISC System/6000	RS/6000
S/370	S/390
SAA	SupportPac
System/390	SystemView
Trouble Ticket	Ultimedia
VisualAge	VTAM
WebConnection	WebExplorer
WIN-OS/2	Workplace
Workplace Shell	400

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Microsoft, Windows, and the Windows 95 logo
are trademarks or registered trademarks of Microsoft Corporation.

Java and HotJava are trademarks of Sun Microsystems, Inc.

ACTION	Prodigy Services Company
Adobe	Adobe Systems, Incorporated
Adobe Photoshop	Adobe Systems, Incorporated
Advantis	Advantis
America Online	America Online, Incorporated
Ameritech	Ameritech, Incorporated
Amiga	Commodore Amiga, Incorporated
Animator	Micro Focus Limited
AppleTalk	Apple Computer, Incorporated
Apple	Apple Computer, Incorporated
Applet	Wilson Window Ware
AT&T	American Telephone and Telegraph Company
Banyan	Banyan Systems, Incorporated
Bristol	Bristol Socket Screw Company
C + +	American Telephone and Telegraph Company, Incorporated
CA	Computer Associates
Cedar	Siemens Company
CheckFree	Checkfree Corporation
Cisco	Cisco Systems, Incorporated

CompuServe	CompuServe Incorporated and H&R Block, Incorporated
Data General	Data General Corporation
DCA	Digital Communications Associates, Incorporated
DCE	The Open Software Foundation
DDS	Sony Corporation
DEC	Digital Equipment Corporation
DECmcc	Digital Equipment Corporation
DECnet	Digital Equipment Corporation
Digital	Digital Equipment Corporation
Discover	Sears, Roebuck and Company
DMS-100	Northern Telecom Limited
DVI	Intel Corporation
ESS	American Telephone and Telegraph Company
Eudora	University of Illinois Board of trustees licensed to QUALCOMM Inc.
Excalibur Technologies	Excalibur Technologies
FrameMaker	Frame Technology, Incorporated
Freelance Graphics	Lotus Development Corporation
Genesis	American Telephone and Telegraph Company
Geneva	Apple Computer, Incorporated
GL	Iris Graphics Library
Gold Disk	Gold Disk Incorporated
Gopher	University of Minnesota
Hayes	Hayes Microcomputer Products, Incorporated
HP/UX	Hewlett-Packard Company
IDNX	Network Equipment Technologies, Incorporated
Indeo	Intel Corporation
Intel	Intel Corporation
Internetwork Packet Exchange	Novell, Incorporated
Interleaf	Interleaf, Incorporated
Internet Phone	VocalTec, Incorporated
IPX	Novell, Incorporated
Kodak	Eastman Kodak Company
LAN WorkPlace	Novell, Incorporated
LMSI	Magnetic Storage International Corporation
Lotus	Lotus Development Corporation
Lotus Notes	Lotus Development Corporation
Lotus 1-2-3	Lotus Development Corporation
Mac OS	Apple Computer, Incorporated
Mac	Apple Computer, Incorporated
Macintosh	Apple Computer, Incorporated
MacOS	Apple Computer, Incorporated
MacTCP	Apple Computer, Incorporated
MCI	MCI Corporation
MCI Mail	MCI Communications Corporation
Meridian	Northern Telecom Limited
Microsoft Windows	Microsoft Corporation
Microsoft	Microsoft Corporation
Microsoft Word	Microsoft Corporation
Milan	Fujitsu PC Corporation
MNP	Microcom Systems, Incorporated
Mosaic	University of Illinois
MOSS	MOSS Systems, Limited
Motif	Open Software Foundation, Incorporated
Motorola	Motorola, Incorporated

MS	Microsoft Corporation
MS-DOS	Microsoft Corporation
MT	NEC Technologies, Incorporated
NAP	Automated Network Management, Incorporated
NCR	NCR Corporation
NCS	Apollo Computer, Incorporated
NCSA Mosaic	University of Illinois at Urbana Champaign
NDIS	3Com Corporation and Microsoft Corporation
Netscape	Netscape Communications Corporation
NetWare	Novell, Incorporated
NeXT	NeXT Computer, Incorporated
NFS	Sun Microsystems Incorporated
Nortel	Northern Telecom
Novell	Novell, Incorporated
ONC	Sun Microsystems, Incorporated
PageMaker	Aldus Corporation
Panasonic	Matsushita Electric Industrial Company, Limited
PC-NFS	Sun Microsystems, Incorporated
Pentium	Intel Corporation
Philips	Philips Consumer Electronics Company
Phoenix	Phoenix Technologies, Limited
Photo CD	Eastman Kodak Company
PostScript	Adobe Systems, Incorporated
PowerPoint	Microsoft Corporation
Prodigy	Prodigy Services Company
QuickTime	Apple Computer, Incorporated
RealAudio	Progressive Networks
Rolodex	Rolodex, Incorporated
SCO	The Santa Cruz Operation, Incorporated
SCSI	Security Control Systems, Incorporated
SmartMasters	Lotus Development Corporation
SmartSuite	Lotus Development Corporation
SmartIcons	Lotus Development Corporation
Solaris	Sun Microsystems, Incorporated
Sony	Sony Corporation
Sprint	Sprint Communications Company
Stac	Stac Electronics
Stacker	Stac Electronics
Sun	Sun Microsystems, Incorporated
SunSoft	Sun Microsystems, Incorporated
Sybase	Sybase Corporation
TI	Texas Instruments Incorporated
Tivoli	Tivoli Systems Inc., an IBM Company
Tivoli Management Framework	Tivoli Systems Inc., an IBM Company
TME 10	Tivoli Systems Inc., an IBM Company
TME	Tivoli Systems Inc., an IBM Company
Toshiba	Toshiba Corporation
Unisys	Unisys Corporation
VINES	Banyan Systems, Incorporated
VIrtual NEtworking Systems	Banyan Systems, Incorporated
VISA	VISA International Services Association
Visual Basic	Microsoft Corporation
VMS	Digital Equipment Corporation
VT	Digital Equipment Corporation
VT100	Digital Equipment Corporation
WebTalk	Quarterdeck Corporation
Wellfleet	Wellfleet Communications, Incorporated
Windows 95	Microsoft Corporation
Win32s	Microsoft Corporation

Wollongong	Wollongong Group
WordPerfect	WordPerfect Corporation
X Window System	Massachusetts Institute of Technology
X-Windows	Massachusetts Institute of Technology
Xerox Network Systems (XNS)	Xerox Corporation
Xerox	Xerox Corporation
386	Intel Corporation
386SX	Intel Corporation
4500	Xerox Corporation
486	Intel Corporation
80386	Intel Corporation

Other trademarks are trademarks of their respective companies.

Appendix I. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

I.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see “How To Get ITSO Redbooks” on page 623.

- *AS/400 Network Management Guide*, GG24-4154
- *Cool Title About the AS/400 and Internet Goes Here*, SG24-4815
- *Network Operations Management - Which Platform? The Practice*, SG24-5015
- *Network Operations Management - Which Platform? The Principle*, SG24-5014
- *SystemView for MVS: Overview and Scenarios*, SG24-4654
- *IBM SystemView for AIX: An Overview*, GG24-2541
- *LAN Network Managers*, SG24-4504
- *NetView for OS/2 as an SNMP Manager*, GG24-4412
- *NetView for AIX V4 Examples*, SG24-4515
- *The Basics of IP Network Design*, SG24-2580
- *Accessing the Internet*, SG24-2597
- *Using the Information Super Highway*, GG24-2499
- *TCP/IP Tutorial and Technical Overview*, GG24-3376
- *Local Area Network Concepts and Products: Routers and Gateways*, SG24-4755
- *IBM 8235 Dial-In Access to LANs Server - Concepts and Experiences*, SG24-4816
- *IBM 2210 Nways Multiprotocol Router Description and Configuration Scenarios*, SG24-4446
- *MPNP V1R3 for IBM 6611*, SG24-4597

I.2 Redbooks on CD-ROMs

Redbooks are also available on CD-ROMs. **Order a subscription** and receive updates 2-4 times a year at significant savings.

CD-ROM Title	Subscription Number	Collection Kit Number
System/390 Redbooks Collection	SBOF-7201	SK2T-2177
Networking and Systems Management Redbooks Collection	SBOF-7370	SK2T-6022
Transaction Processing and Data Management Redbook	SBOF-7240	SK2T-8038
AS/400 Redbooks Collection	SBOF-7270	SK2T-2849
RISC System/6000 Redbooks Collection (HTML, BkMgr)	SBOF-7230	SK2T-8040
RISC System/6000 Redbooks Collection (PostScript)	SBOF-7205	SK2T-8041
Application Development Redbooks Collection	SBOF-7290	SK2T-8037
Personal Systems Redbooks Collection	SBOF-7250	SK2T-8042

© Copyright IBM Corp. 1996

I.3 Other Publications

These publications are also relevant as further information sources.

- *Teach Yourself Java in 21 Days*, ISBN 1-57521-030-4
- *Java in a Nutshell*, ISBN 1-56592-183-6
- *Hooked on Java*, ISBN 0-201-48837-x

How To Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, CD-ROMs, workshops, and residencies. A form for ordering books and CD-ROMs is also provided.

This information was current at the time of publication, but is continually subject to change. The latest information may be found at URL http://www.redbooks.ibm.com.

How IBM Employees Can Get ITSO Redbooks

Employees may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **PUBORDER** — to order hardcopies in United States
- **GOPHER link to the Internet** - type GOPHER.WTSCPOK.ITSO.IBM.COM
- **Tools disks**

 To get LIST3820s of redbooks, type one of the following commands:

  ```
  TOOLS SENDTO EHONE4 TOOLS2 REDPRINT GET SG24xxxx PACKAGE
  TOOLS SENDTO CANVM2 TOOLS REDPRINT GET SG24xxxx PACKAGE (Canadian users only)
  ```

 To get lists of redbooks:

  ```
  TOOLS SENDTO WTSCPOK TOOLS REDBOOKS GET REDBOOKS CATALOG
  TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET ITSOCAT TXT
  TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET LISTSERV PACKAGE
  ```

 To register for information on workshops, residencies, and redbooks:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ITSOREGI 1996
  ```

 For a list of product area specialists in the ITSO:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ORGCARD PACKAGE
  ```

- **Redbooks Home Page on the World Wide Web**

 http://w3.itso.ibm.com/redbooks

- **IBM Direct Publications Catalog on the World Wide Web**

 http://www.elink.ibmlink.ibm.com/pbl/pbl

 IBM employees may obtain LIST3820s of redbooks from this page.

- **REDBOOKS category on INEWS**
- **Online** — send orders to: USIB6FPL at IBMMAIL or DKIBMBSH at IBMMAIL
- **Internet Listserver**

 With an Internet E-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an E-mail note to announce@webster.ibmlink.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank). A category form and detailed instructions will be sent to you.

© Copyright IBM Corp. 1996

How Customers Can Get ITSO Redbooks

Customers may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Online Orders** (Do not send credit card information over the Internet) — send orders to:

	IBMMAIL	**Internet**
In United States:	usib6fpl at ibmmail	usib6fpl@ibmmail.com
In Canada:	caibmbkz at ibmmail	lmannix@vnet.ibm.com
Outside North America:	dkibmbsh at ibmmail	bookshop@dk.ibm.com

- **Telephone orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	(long distance charges apply)
(+45) 4810-1320 - Danish	(+45) 4810-1020 - German
(+45) 4810-1420 - Dutch	(+45) 4810-1620 - Italian
(+45) 4810-1540 - English	(+45) 4810-1270 - Norwegian
(+45) 4810-1670 - Finnish	(+45) 4810-1120 - Spanish
(+45) 4810-1220 - French	(+45) 4810-1170 - Swedish

- **Mail Orders** — send orders to:

IBM Publications	IBM Publications	IBM Direct Services
Publications Customer Support	144-4th Avenue, S.W.	Sortemosevej 21
P.O. Box 29570	Calgary, Alberta T2P 3N5	DK-3450 Allerød
Raleigh, NC 27626-0570	Canada	Denmark
USA		

- **Fax** — send orders to:

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	(+45) 48 14 2207 (long distance charge)

- **1-800-IBM-4FAX (United States)** or **(+1) 415 855 43 29 (Outside USA)** — ask for:

 Index # 4421 Abstracts of new redbooks
 Index # 4422 IBM redbooks
 Index # 4420 Redbooks for last six months

- **Direct Services** - send note to softwareshop@vnet.ibm.com
- **On the World Wide Web**

Redbooks Home Page	http://www.redbooks.ibm.com
IBM Direct Publications Catalog	http://www.elink.ibmlink.ibm.com/pbl/pbl

- **Internet Listserver**

 With an Internet E-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an E-mail note to announce@webster.ibmlink.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank).

IBM Redbook Order Form

Please send me the following:

Title	Order Number	Quantity

- **Please put me on the mailing list for updated versions of the IBM Redbook Catalog.**

First name Last name

Company

Address

City Postal code Country

Telephone number Telefax number VAT number

- Invoice to customer number

- Credit card number

Credit card expiration date Card issued to Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

DO NOT SEND CREDIT CARD INFORMATION OVER THE INTERNET.

Index

Special Characters

$include definition 467
$origin definition 467, 468, 470, 471
%SQL subsections on DB2 gateway 305

Numerics

1.5 Mbps 17
10BaseT 18
128 kbps 18
2210 Configuration Program
 Configuration Window 98
 hardware requirements 98
 Navigation Window 98
 software requirements 98
2210 Nways Multiprotocol Router
 indicators 91
 local access 93
 models 89
 remote access 93
 reset button 92
 supported networks 93
32-bit 458
44.6Mb/s 17
5270 to HTML gateway
 AS/400 5250-to-HTML server 329
 example of legacy DDS 333
 logon exit program API 330
 new HTML DDS keyword 335
56 kbps 17
64 kbps 17
6611 Router
 Model 120 124
 Model 125 125
 Model 145 126
 Model 175 126
8235
 8235 hardware description 71
 activity logger 88
 AppleTalk 83
 bridging 77
 code structure 74
 communication options 76
 description 21
 DIALs client software 22
 IP traffic 79
 IPX traffic 81
 LAN connection 72
 LED, network and port status 71
 LED, power status 71
 management facility 86, 88
 microcode 75
 models summary 75
 NetBIOS and 802.2 77
8235 *(continued)*
 security 86
 supported protocols 77
 system components 21
 technical description 21
8235 BRI module 33
8235 Management Facility 21, 46
 8235 User List 63
 adding devices to an IP device list file 49
 Async Serial Port Configuration dialog box 56
 auto-download 48
 clear and download 48
 Devices menu 48
 Digital Pathways 67
 Discover Devices 48
 downloading of VROM and image files to the 8235 48
 Ethernet 48
 General Configuration page 52
 hardware requirements 46
 IBM 8235 Program Group 47
 Internal Modem Module Port Configuration dialog box 53
 Internal User List 62
 Internet applications 47
 IP Addresses Configuration page 59
 IP Device window 50
 IP General Configuration page 58
 IP Static Routes Configuration page 61
 Management Protocols page 50
 NetWare Bindery 63
 Ports Configuration page 53
 Ports:Phone Numbers Configuration page 56
 Preferences window 50
 Radius 66
 Routing Table window 70
 SecurID 67
 Security Configuration page 62
 service providers 47
 SNMP Configuration page 69
 software requirements 46
 TACACS 64
 TACACS Plus 65
 Token-Ring 48
 Virtual Connections Configuration page 57
8235 Management Facility for Windows 42
8235 Model I40 DIALs Switch 55

A

absolute name 464
Abuse of privilege 362
access control system 105
Access points 346

© Copyright IBM Corp. 1996

Accessing remote DB2 databases 299
Activity 344
Adaptive Rate-Based (ARB) 153
Adaptive Source-Routing Transparent Bridge (ASRT) 118
additional routing 90
Address Resolution Protocol (ARP) 79, 80
address tables 459
address to name translation 476
admin-bin 214
Administration 164
administration form 507
administrative check 505
administrative forms 503
ADSTAR Distributed Storage Manager/2 (ADSM/2) 574
age out timer 147
ain names 356
AIX 159, 160, 162, 163
AIX Platform
AIX NetView Service Point 406
current product releases 405
Distributed 403
LAN Management Utilities/6000 405
LAN Remote Monitor for AIX 405
LAN Workgroup 403
main product 405
managing IP networks 404
Netview for AIX 405
Network Manager for AIX 405
Nways BroadBand Switch Manager 406
Nways Campus Manager ATM for AIX 405
Nways Campus Manager for AIX 406
Nways Campus Manager LAN for AIX 406
overview 405
Positioning 403
Router and Bridge Manager/6000 405
skill requirement 404
SNA Manager/6000 405
Systems Monitor for AIX 406
Telecommunications Management Network Product 406
Trouble Ticket for AIX 406
AIX Transmission Network Manager/6000 567
alias name 469
ALIGN 183
ALIGN variable 301
all-routes broadcast 133
Analysis 341
Anchor 182
Animations with .GIF files 260
anonymous 513
anonymous FTP 512, 514, 558
anonymous FTP area 513
anonymous FTP area public directories 514
anonymous FTP site 99
anonymous users 515
ANR (Automatic Network Routing) 153
Antennas 13
AON/MVS 571
Apache HTTP Server 232
API (Application Programming Interface) 44
APNIC 480
applet 249
AppleTalk 94, 130
AppleTalk ARA 2.0 83
AppleTalk broadcast packets 84
AppleTalk Remote Access Protocol (ARAP) 83
Applets viewers 251
Application level 355
Application Programming Interface (API) 44
ARA 2.0 41
ARA routers 84
ARAP (AppleTalk Remote Access Protocol) 83
ARB (Adaptive Rate-Based) 153
Archie 358
archive format 96
ARP 356
ARP (Address Resolution Protocol) 79, 80, 103
Arpanet 559
Arrays object 237
AS boundary routing 107
AS/400 1, 161
AS/400 FSIOP 10
AS/400 native applications on the web 10
AS/400 Notes support 10
AS/400 POP3 implementation 10
AS/400 security 11
ASCII
console 112
emulator 93
terminal 93, 95
ASRT (Adaptive Source-Routing Transparent Bridge) 118
assistance 554
Asynchronous Request/Response Protocol 394
ATN 18
audio 207
Audio .aif,.aiff and .aifc 283
Audio .au and .snd 283
Audio .mod format 284
Audio .wav format 283
Audio formats 283
AudioClip object in Java 263
Auditing 369
AUI(Thick Ethernet) 73
AUTH_TYPE 219
authentication 107, 362
Authorization 362
Automatic Network Routing (ANR) 153
automatic reconnection 31
Average web response size 2
Average web transaction size 2
AVI 294

B

B channels 33
Backbone 441
Backup 352
backup switched line 120
BAN (Boundary Access Node) 154
Bandwidth 1, 118
bandwidth reservation (BRS) 117
Banyan VINES Control Protocol (BVCP) 102
Banyan Virtual Networking System Protocol (VINES) 102
Banyan Virtual Networking Systems (VINES) 129
Basic Primary Rate Service 19
basic training 557
Bastion host 356
BGP (Border Gateway Protocol) 131
bibliography 621
bit permissions 521
BNC(Thin Ethernet) 73
BODY 176
BookMaster 208
boot configuration database 94
boot files 94
boot processes 94
Boot PROM 74
Boot Protocol (BOOTP) 80
BOOTP (boot protocol) 80, 95
 client 111
 forwarder 111
 relay agent 111
 server 111
Bootstrap Protocol 104
Border Gateway Protocol (BGP) 131
Boundary Access Node (BAN) 154
bracket 176
bracket, cable management 127
BRI module 76
Bridging 77
 LAN-to-LAN 100
 LAN-to-WAN 100
bridging switch 102
broadcast address type 106
browsing capabilities 204
BRS (bandwidth reservation) 101, 117
BSD Compress-LZW 101
building a content service 484
Buildings 16
Business 422
Business management 387
Business transformation projects 553
BVCP (Banyan VINES Control Protocol) 102

C

Cable Company 18
cable management bracket 127
Cable modems 17
cache 147
caching-only name server 473
Calculating HTTP operations 5
Calling classes to create new objects 255
Campus 340
Canvas object 269
Capacity planning example 5
Capacity requirements 5
Care 340
catalog outsourcing 536
catch 255
CCITT 389
CCL (Command Control Languages) 87
CD-ROM 294
Centralized, Definition 402
CERT 342
CGI 512
 C language 219
 Decoding the input from a form 217
 GET Method 216
 How to create a. 215
 Interesting places 228
 Meaning 214
 PERL 219
 POST Method 216
 REXX 219
 Transferences methods of a 215, 216
CGI Programs 509, 512
CGI script 192
cgi-bin 214
Challenge 362
Change management 387
channel aggregation 45
check boxes 268
Checksums 366
child nodes 463
Choice object 269
Chroot 362
CICS 11
class names 118
class resource record 465
client event logging 43
Client software 155, 158
client's business strategy 557
client/server 110
cmip 389
CNAME resource record 469
Color object 258
COM21 18
Command Control Languages (CCL) 87
Common layouts 273
Common sense 347
Communications programs 341
community 112
Compact discs 287
compatibility mode bridging 135
Compuserve GIF 283

Computer users 340
Computers 341
Confidentiality 364
configuration 165
Configuration management 387
 configuration report server (CRS) 101, 145
 Configuration Window 98
Configuring a server 6
 connect application 22
connection 456
 connection file 23
 Connection File Wizard 43
connection service 456
Connection speed 1
connectionless protocol 393
 connectivity, multiprotocol 129
Consideration 340
 console port 94
 consultancy 553
 Consulting Group 553
 Consulting Service Lines 553
Consulting Services 157, 553
content hosting 455
content provider for infoMarket 525
content services 451, 557
content services network 457
Content type 2
 CONTENT_LENGTH 216
CONTENT_TYPE 219
Controls 347
convert 208
coordinates of the image 187
Copyright 343
Corporate Dial Services
 Eudora Light 432
 Local Dial 432
 Netscape 432
 POP3 432
 SLIP 432
 User Network Authentication 433
corporate intranets 556
Cost 341
Cost of protecting 339
country domain 465
CRC 366
Creation 340
Critical 340
 CRS (configuration report server) 101, 145
Crypt 365
Cryptolope 526
Cryptosealing 366
CyberCash 377

D

DAC 294
DARPA 388
 data compression 32
Data driven 362
 Data Link Connection (DLC) 134
 Data Link Connection Identifier (DLCI) 134
data link switching (DLSw) 114, 123, 130, 145, 147
 database 138
DATABASE variable 301
 database, Ethernet 141
 database, Token-Ring 141
DatagLANce 574
 Date class 251
DB_CASE variable 301
DB/2 11
 DB2 bindfile 301
 DB2 gateway initialization file 301
DB2WWW availability 299
 DB2WWW plattaforms 299
DDN Network Information Center (NIC) 459
Decisions 340
 DECnet 129
 DECnet IV 102, 103
 DECnet V 102
default directory 508
default gateway 105, 110
default name server 476
default router 110
default subnet gateway 110
default subnetwork gateway 105
Defense Information Systems Agency (DISA) 459
DEFINE section and keyword for DB2WWW 302
definition list 178
 Deflate - LZ77 100
 Delta technology 45
Denial 342
Denial of service 358
 deny mode 137
Department of Defense (DoD) 459
 Dependent LU Requester (DLUR) 154
DES 365
 designated ring 135
 Destructor 243
 dial back 42
 dial-in access 21
 dial-in channel aggregation 43
dial-on-demand 121
Dial-up 350
Dial-Up Services
 Connection Scripts 429
 Eudora Light 428
 IBM Internet Connection Access Kit 428
 IBM Internet Connection for Windows 428
 IBM Internet Dialer 429
 IBM OS/2 Warp Internet Access 428
 IBM WebExplorer for OS/2 428
 IBM WebExplorer Mosaic 428
 IP Address 429
 Local Dial 430
 Netscape 428
 OS/2 Warp 431

Dial-Up Services *(continued)*
OS/2 Warp Connect 431
POP3 430
Services 430
SLIP 429
Trumpet Winsock 428
Windows 3.1 431
Windows 95 431
DIALs Client 22
Advanced ISDN Settings dialog box 33
Advanced Settings dialog box 31
communications ports 29
Connection File Options dialog box 26
DIALs Connect/2 window 24
Edit Modem Configuration dialog box 30
modem settings 29
Modem Setup dialog box 30
NDIS (Network Driver Interface) 25
Network Driver Interface (NDIS) 25
ODI (Open Data-Link Interface) 25
Open Data-Link Interface (ODI) 25
port settings 29
Port Setup dialog box 29
roaming dial-back 28
DIALs Client/2 22
DIALs Connect/2 22, 23
DIFF 349
Different security 340
Digital movie formats 286
Digital phone-line 17
Digital video file formats 287
Digital video hardware requirements 286
Digital video players 288
Digital video software requirements 287
Direct Satellite Broadcast 287
directory list 179
DISA 459
Disclosure 342, 521
Distance-Vector Multicast Routing Protocol (DVMRP) 110
Distributed Console Access Facility (DCAF) 574
distributed database 460
Distributed, Definition 402
Dividing daemons 4
DLC (Data Link Connection) 134
DLC termination 114
DLCI (Data Link Connection Identifier) 134
DLCI (PVC) 108
DLSw (data link switching) 104, 114, 123, 130, 145, 147
DLSw connections 116
DLSw partner 116
DLUR (Dependent LU Requester) 154
DMI 400
Component Interface 401
DCE/RPC 401
Desktop Management Taskforce 400
elements of 401
DMI *(continued)*
Management Interface 401
MIF 401
ONC/RPC 401
Service Provider 401
TI/RPC 401
DNS 356, 459, 460, 463, 469
alias 489
dependencies 480
design 479
DNS 459, 460, 463, 469
mail services 479
mapping of addresses 461, 469
mapping of domain names 469
mapping of names 480
master file 467, 469
master file control entries 467
messages 473
name resolution 480
name resolution services 481
name resource record 465
name space structure 461
name systems 459
name to address translation 475, 480
namespace 460, 469, 472
namespace explosion 460
namespace structure 481
queries 473
resource names 469
security 482
zone 471
DNS alias 489
DNS dependencies 480
DNS design 479
DNS mail services 479
DNS master file 467, 469
DNS messages 473
DNS queries 473
DNS resource names 469
DNS security 482
DNS zone 471
do...while 247
DoD 459
Domain Name Services 436
Domino 167, 168
Domino applications 170
Domino Benefits 168
Domino Features 169
DOS drivers 22
Double buffering. 258
downloadable files 514
Downstream 17
DRAM 89
DSP 13
Dual homed gateway 363
dual mode bridging 142
Dual Speed 294

DUMP 352
DVMRP 103
DVMRP (Distance-Vector Multicast Routing Protocol) 110
Dynamic content 2

E

e-mail 482, 557
E1 423
EasyStart 100, 154
ECPA 344
Educating 351
EGP 110
EGP (Exterior Gateway Protocol) 110, 131
Electronic Mail 424
Electronic Purchasing Service (EPS) 534
Eletronic commerce 374
ELS (Event Logging System) 97
emerging technologies 554
emulator 93
Encryption 364
end node 83
Enhanced Priority Queueing 154
enterprise′s primary processes 555
enterprise-specific traps 113
enterprise-wide information systems 554
EPS
 electronic catalog 534
 features 535
 procurement process 534
equals() 263
error conditions 495
Error handling for DB2 305
ETC\HOSTS file 479
Ethernet 2, 459
 database 141
event logging system (ELS) 97
event number 97
Exchange ID (XID) 115
executable memory 95
express installation 43
extended characters 198
Exterior Gateway Protocol (EGP) 110, 131

F

facilities 169
FAQ about capacity planning 6
Fast Ethernet 2
FAX 354
FDDI 2
Features 161, 166
File Transfer Protocol 512
fill-pattern 105
filtering database 138
filters 110
final, Java keyword 268
finalize() 243
Firewall 353
firewall mail server 484
firewall name server 483
Firewall software 156
firewalls 482
firewalls and electronic mail 483
Fix 340, 352
flash memory 89
flat name space 460
flat namespace 460
Floating Virtual Connections (FVC) 44
Font selection in Java 249
FontMetrics class 257
for() 247
Form 188
Forms 188, 341
four-port serial adapter 128
FQDN 463
FR Boundary Access Node (BAN) 154
Frame class 275
frame relay interface 108, 109
frame types 82
FrameMaker 212, 214
Frequency 12
FTP 357, 452, 476, 512, 557
 anonymous 513
 anonymous FTP 512, 514
 anonymous FTP area 513
 anonymous FTP area public directories 514
 anonymous users 515
 File Transfer Protocol 512
 FTP 452, 476, 512
 FTP access 508
 incoming area 515
 incoming directory 515
FTP access 508
full resolver 476
full resolver example 477
Full-color video 19, 279
Full-motion video 17
fully meshed 109
fully qualified domain name (FQDN) 463
FVC (Floating Virtual Connections) 44

G

Game playing 350
GATEWAY_INTERFACE 218
General Instrument 18
General MIDI standard 285
general traps 113
generic TLDs 465
geographical domain 465
GET Method 222
GET method 216
getAppletInfo 276
getAudioClip() 263

getCodeBase() 262
getImage() 260, 263
getParameter() method 252
getstats 517
getstats program reports 517
GIF 204, 263, 281
GIF format 204
GIF Frames 281
GIF image 186, 206
GIF limitations 283
GIF logical screen area 281
GIF, benefit to use 282
GIF87a 281
GIF89a 281
GMFHS 569
Goals 340
good performance 1
Gopher 358
Gopher servers 558
graphic 173
Graphic Web browsers 204
Graphics Interchange Format 204
Gray-scale video 279
group 97
GTE 377
GWCON 96

H

hacker 2
handleEvent (Event e) 266
Hardware 341, 455
Hardware and software combination 2
hardware flow control 31
Hardware/Software Requirements 162
Hayes-Compatible Modem 431, 434
HEAD 176
heading levels 178
heterogeneous network 388
hex numbering system 457
hierarchical name space 460, 472
hierarchical tree structure 461
High Performance Routing (HPR) 153
high-bandwidth Internet link 557
High-definition television 287
High-speed connection 14
Hijacking 363
home page 501, 557
hop count 136, 137
host name 475
hostname resolution 497
HOSTS.LOCAL file 479
HOSTS.TXT file 472
How to convert OS/400 screens to HTML
See 5270 to HTML gateway
HP 18
HP-UX 163
HPR (High Performance Routing) 153
HTML 171, 172, 175, 208, 212, 213, 214
conversion from BookMaster to HTML 208
conversion from FrameMaker to HTML 212
conversion from interleaf to HTML 213
conversions 214
converters 214
description 175
documents 175
editors 201, 202
Form 190
HTML+ 193
HTML+ complex tables support 196
HTML+ large documents 195
HTML+ math tags 197
HTML2.0 document structure 176
HTML2.0 syntax 176
HTML3.0 193
HyperText Markup Language 175, 208
language 175
Main Elements 176
HTML converters 214
HTML documents 175
HTML editors 201, 202
HTML Form 190
HTML language 175
HTML Main Elements 176
HTML+ 193, 195, 196, 197
HTML2.0 Document Structure 176
HTML2.0 Syntax 176
HTML3.0 193
HTTP APIs 232
HTTP port number 491
HTTP_ACCEPT 219
HTTP_USER_AGENT 219
Hugues 18
hyperlink 182
HyperText Markup Language 175, 208

I

I/O 3
I/T Consulting Services
Application Development Effectiveness 554
Business Recovery 554
I/T Planning 554
Information Systems Management 554
Networking 554
I/T solutions 555
IAB 459, 559
IAB, function 561
IANA 480, 481, 560
IBD (Integrated Boot Device) 94, 113
IBM 18, 426
IBM 2210 Nways 425, 435
IBM 6611 425
IBM 6611's family 124
IBM AS/400 10
IBM Connection Server Family 161

IBM Content Services 453
IBM Dial-Up for TCP/IP 34
 Add Entries Window 35
 Connect Info window 37
 IBM Dial-Up for TCP/IP window 34
 Login Info window 36
 login sequence 35
 Mail Server information 38
 Modem Info window 39
 Modify Entries window 36
 Network Dialer 34
 PPP (Point-to-Point Protocol) 34
 Server Info window 38
 SLIP (Serial Line Internet Protocol) 34
IBM Global Network 453
 Advantis 426, 432
 Customer Support 428
 IBM Information Network 426
 IGN 434
 Joint Ventures 426
 Points of Presence 428
 Prodigy 426
IBM Global Network Content Services 156
IBM Global Network Internet Connection 156
IBM infoMarket 523
 content provider 525
 controlled use of information 526
 Cryptolope Helper 529
 IBM infoMarket News Ticker 530
 information distribution 527
 Plug-N-Publish Toolkits 531
 Rights Management Architecture 577
 search by source 524
 search feature 524
 secure container architecture 526
 secure payment 527
 undetected alteration 526
IBM InfoMarket Service 156
IBM infoSage 531
 Archive Search 533
 Content Resources 581
 IBM Profile Editor Screenshots 584
 information delivery system 531
 Links 533
 newsletter 531
 Profile Editor 532
 Special Editions 533
 Stock Tracker 533
 Top Stories 533
 topics 532
IBM Internet Connection Access Services 419
IBM Internet connection server family 162
IBM offers 155
IBM PC Server 6
IBM PowerPC 425
IBM RISC/6000 425
IBM RS/6000 8, 429
IBM S/390 11
IBM Servers 6
IBM SystemView for OS/2 573
IBM WaveRunner digital modem 31
IBM WebExplorer 41
ICMP 103
ICMP (Internet Control Message Protocol) 80
id property of the Event object 266
if...else 245
ifconfig command 485
IGMP 109
IGN 453
iKP 377
Image 183, 186
 converters 207
 format 206
 image 183, 186
 Image Map 186, 188
 imagemap 187
 imagemap program 186
 images 183, 204
 ISMAP 186
image converters 207
image format 206
Image Map 186, 188
imagemap 187
imagemap program 186
Images 2, 183, 204
Images Formats on Java 263
Implementation 340
implements Runnable 254
IMS 11
in-addr.arpa 469, 470, 476
In-house applications 4
InARP 102
inbound filters 136
incoming area 515
incoming directory 515
index.html 499
Industry Specializations
 Cross Industry 556
 Distribution 555
 Finance 555
 Government 555
 Healthcare 555
 Higher Education 555
 Insurance 555
 Manufacturing 555
 Petroleum 556
 Telecommunications and Media 556
 Travel 556
inetd 469
Information Gateways 156
Infrastructure investment 5
INGW 429
init() method declaration 252
INPUT 188

input option invoking DB2WWW 302
INRB 481
installation 487
Installation checklist on DB2WWW 300
Integrated Boot Device (IBD) 94, 113
Integrated Solution Offering (ISO) 557
Integration Services
 Life Cycle 554
 Rapid Solutions 555
 Redevelopment 555
 Systems Integration 555
Intel 1
Interactive Marketing Service 536
Interface 6
Interlacing 281
Interleaf 213, 214
intermediate nodes 146
intermediate-sized businesses 555
internal IP address 104
internal name server 483
Internet 422
Internet applications 47
Internet Architecture Board (IAB) 459
Internet business 556
Internet Configuration Control Board 559
Internet Connection 556
 IBM Internet connection server family 162
 Internet Connection for OS/2 158
 Internet Connection for Windows 158
 Internet Connection Secure Server 160
 Internet Connection Server 159, 160, 163
Internet Connection for OS/2 158
Internet Connection for Windows 158
Internet Connection Secure Server 160
Internet Connection Server 159, 160, 163
Internet Control Message Protocol 103
Internet Control Message Protocol (ICMP) 80
Internet DNS Names Review Board (INRB) 481
Internet Engineering Steering Group 559
Internet Engineering Task Force 559
Internet Host Table 460
Internet Implementation Integrated Solution Offering 557
Internet Network Information Center 459
Internet offerings 155
Internet Package Exchange (IPX) 23
Internet Packet Exchange (IPX) 82, 129
Internet POWERsolution 163
Internet Protocol (IP) 80, 129
Internet protocol state definitions 562
Internet protocol status definitions 562
Internet Registry (IR) 480
Internet Research Steering Group 559
Internet Research Task Force 559
Internet servers 156, 163
Internet Services 423
Internet solution 557
Internet solutions 155
Internet strategy 556
InterNIC 433, 436, 459, 480
 domain requirements 481
 Internet Network Information Center 459
 InterNIC 459, 480
 registering a domain 480
 registration 459, 481
 registration templates 481
InterNIC domain requirements 481
InterNIC registration 459
InterNIC registration templates 481
InterNotes News 166, 167
InterNotes Web Publisher 164, 165, 166
InterNotes Web Publisher 4.0 165
intranet 1
intranet solution 557
Invoking the DB2 gateway 302
IP 23, 27
IP (Internet Protocol) 80, 118, 124, 129
 access control 105
 broadcast format 105
 encapsulation 114
 filters 110
 internal address 104
 routing 103, 131
 tunneling 110
IP address 27, 439, 457, 458, 459, 461, 464, 469, 480, 490, 496
IP address alias 485
IP download 45
IP protocol 459
IPGATEWAY 85
IPX (Internet Package Exchange) 23, 27
IPX (Internet Packet Exchange) 82, 94, 118, 129
IPX router function 81
IR 480
ISDN 1, 19, 89, 93, 101, 423
ISDN service example 19
ISDN terminal adapter 30
ISMAP 186
ISO 389
ISO 3166 465
Isolated 340
ISSC Consulting and Services 557
ISSC Cross Industry 557
iterative mode 474

J

Java 8
Java Abstract Window Toolkit
 Adding components 268
 Buttons 268
 Components 268
 Constructors 269
 Labels 268
 Methods 270

Java Availability 236
Java books 275
Java Compiler characteristics 237
Java operators 239
Java Virtual Machine 235
Java′s API 270
JavaScript 233, 277
JGEG 263
Joint Photographic Expert Group 205
JPEG 205, 279
JPEG compression 279
Juggling Virtual Connections (JVC) 44
JVC (Juggling Virtual Connections) 44

K

Kerberos 367
Key element 340
Kinetics Internet Protocol (KIP) 85
KIP (Kinetics Internet Protocol) 85

L

LAA (locally administered address) 115, 152
LAN 1
LAN Bridge Server (LBS) 145
LAN Bridging Protocol 135, 144
LAN Management Utilities/6000 565
LAN Network Manager for AIX 566
LAN Network Manager for OS/2 (LNM) 573
LAN Remote Monitor for AIX 566
LAN Remote Monitor for Windows 575
LAN Reporting Mechanism (LRM) 145
LAN Workgroup, Definition 402
LAN-to-LAN bridging 100
LAN-to-WAN bridging 100
LAN/WAN combinations 128
LANConnect applets 45
large documents 195
Large-volume transactions 11
Laws 340
layer address, network 131
LBS (LAN Bridge Server) 145
Leased Line 427
Leased line connections 16
Leased Line Services
 Customer Support: 436
 Direct Leased Line Internet Access 435
 LAN Internetworking 435
 Leased Line 434
 Network Management: 436
 TCP/IP Connectivity 435
 TCPGATE2 436
Leased lines 1
Least privilege 363
Levels of responsibility 340
LIG 429
limiting access 520
link 456
link speed 456
Linux 3
List servers 558
Lists 178
LLC 23, 27
LLC (Logical Link Control) 41
LLC SAP filters 77
LLC type 2 147
local bridges 132
local data link switching 115, 147
local-wire 105
locally administered address (LAA) 115, 152
log directory 516
Logging 363
logging level 97
Logical Link Control (LLC) 41
LOGIN variable 301
Lossy comrpession 280
Lotus Domino WebServer
 applications 170
 Benefits 168
 Domino 167, 168
 facilities 169
 Features 169
 Requirements 168
 Web site 169
Lotus InterNotes 156, 163
 configuration 165
 InterNotes News 166, 167
 InterNotes Web Publisher 164, 165, 166
 InterNotes Web Publisher 4.0 165
 Requirements 165, 166
Lotus Notes 168
Lotus Word Pro
 requirements 173
 Word Pro 170, 172, 173
 word processors 171
LRM (LAN Reporting Mechanism) 145
LS 349
LZW compressed images 281

M

MAC 371
MAC address 459
MAC Filtering (MCF) 102
MacTCP 41
Magnetic media 341
Mail
 forwarding 483
 mail server 479, 483
 mailbox name 479
 routing 479
 support 479
mail forwarding 483
mail routing 479
mail server 479, 483

mail support 479
mailbox name 479
MAINT 351
management 519
management facility 45
management information base (MIB) 113
management software 519
map 183, 186
mapping of addresses 461, 469
mapping of domain names 469
mapping of names 480
Master Card 377
master file 467, 469
master file control entries 467
mastering 43
mathematical 197
maximum number of hops 106, 112
MCF (MAC Filtering) 102
MDC 366
Media Access Control (MAC) 459
Media Access Control Bridges (802.1D) 138
menu list 180
MIB 399
MIB (management information base) 113
MIB, Enterprise-specific 399
MIB, Experimental 399
MIB, Standard 399
Microwave radio 13
MIDI 284
MIDI channels 285
MIDI device 284
MIDI mapper 285
MIDI sequencer 285
Military 340
MIT 367
MLP (Multilink PPP protocol) 42
Model 120 enhancements 124
Model of security 339
model upgrade 127
monitor 520
MONITOR process 97
monitoring system 94
Monitoring tools 348
MOSPF (Multicast OSPF) 109
most recent router 83
Mountains 15
Mouse event handler methods 267
mouseDown (Event e, int x, int y) 267
mouseDrag(Event e, int x, int y) 267
mouseUp(Event e, intx, int y) 267
MPEG 287
MPEG-2 287
MPNP (Multiprotocol Network Program) 124
MQSeries 11
MQSeries Gateway
amqwget0.cpp 323
amqwput0.cpp 323
cache C++ files 324
MQSeries Gateway *(continued)*
CGI to MQSeries convertion 311
CGIPart 323
CGIPartSet 323
Configuration 312
ConfigurationPart 323
ConfigurationSet 323
Default setting 314
Directory structure 322
Files 310
formRequest 322
formResponse 322
Gateway.Replay.Queue 311
hq,sgda C++ files 324
hqmsfmd C++ files 324
hqmsglist C++ files 324
HTMLPage 323
MQGate 311
MQGate.ini 312
MQGateway 322
MQHost.cpp 323
MQIGwQueue 312
MQIGwQueueManager 312
MQIGwReplyQueue 312
MQIGwWaitInterval 312
MQQueueB 313
MQQueueB.cpp 323
queuepro C++ files 323
queuescan C++ files 323
StatusPage 323
Tested scenarios 310
URLDecoder 323
MRNS (Multiprotocol Routing Network Services) 94
MRNS Configuration Program 96
MRNS user interface 97
MSM 569
multi-interface serial adapter 129
multi-user host 482
Multicast OSPF (MOSPF) 109
multihomed hosts 482
Multilink 33
Multilink PPP protocol (MLP) 42
Multimedia glossary 293
Multimedia web content 2
Multiple GIF images 281
Multiple home-pages 6
multiple IP addresses 485
Multiple Retrieve Function 154
Multiple strategies 348
Multiprocessing with AIX 3
Multiprocessing with OS/2 4
Multiprocessors 3
multiprotocol connectivity 129
multiprotocol network 114
Multiprotocol Network Program (MPNP) 124
Multiprotocol Routing Network Services (MRNS) 94
multivendor 556

MVS Platform
APPNTAM 409
Centralized 407
current product releases 409
Distributed 407
GMFHS 409
Information Management 410
intermediate manager 408
LAN Workgroup 407
MSM 409
NetView AutoBridge/MVS 410
NetView for MVS 409
NetView Remote Operations Manager MVS 409
NGMF 409
NNP/2 410
non-SNA environment 407
NPM 410
OSI/CS 410
overview 408
positioning 406
RODM 409
service point manager 408
Six2View 410
skill requirement 407
SNA network 407
MVS/ESA 160
MX resource record 466, 469, 479

N

N_column-name 306
N1...Ni Variables 306
Name Binding Protocol (NBP) 85
name resolution 480
name resolution services 481
name resource record 465
name space structure 461
name systems 459
name to address translation 475, 480
namespace 460, 469, 472
namespace explosion 460
namespace structure 481
native mode bridging 134
Navigation Window 98
NBMA (non-broadcast multiaccess) 107
NBP (Name Binding Protocol) 85
NDIS 22
Net.Commerce 537
Daemon 539
Director 539
electronic store 538
Lotus Payment Switch 540
merchant 537
Store Administrator 539
Store Creator 538
Store Manager 538
Template Editor 539
NetBEUI 23, 27
NetBEUI (NetBIOS Extended User Interface) 41
NetBIOS 124
data link switching 153
devices 153
frame size reduction 152
names 147
NetBIOS Extended User Interface (NetBEUI) 41
NetFinity Manager for Windows 575
Netscape 8
Netscape browser for VRML 230
Netscape HTTP server 487
Netscape Server 163, 487
NetSP 373
NetView
AIX Service Point 565
APPNTAM 570
AutoBridge/MVS 570
Bridge 570
Distribution Manager/2 574
DM Easy Preparer for OS/2 574
for AIX 565
for MVS 568
for OS/2 573
for Windows 575
LAN Management Utilities 573
NetView/PC 569
Network Planner/2 572
Performance Monitor 571
Remote Operations Agent/400 572
Remote Operations Manager MVS 572
Service Point for AIX 569
TMN Support Facility for AIX 567
Netware Bindery 42, 86
network 457
network administrator 23, 27
network based applications 556
network computing 523
Network Door/2 (NetDoor) 574
Network Information Center 560
network layer address 131
Network level 355
network management 519
Network Management Products 565
Network Management, Overview 402
Network managers 341
network protocols 23, 27
network security 558
Network Services 156
Networked Applications 523
networking 456
NETX 42
new adapters 125, 126, 128
new port driver 43
News servers 558
NGMF 568
NIC 17, 459, 464, 472, 560
NLIST Variable 306

node 461, 463, 464, 465, 471
non-broadcast 107
non-broadcast multiaccess (NBMA) 107
non-recursive mode 474
non-secure network 482
non-translational bridge 143
non-volatile RAM (NVRAM) 96
Note 157
Notes 168
Notes database 164
Novell NetWare 7
Novell UnixWare 7
NS resource record 473, 475
NSF 388
NUM_COLUMNS Variable 306
number of processes 494
NVRAM (non-volatile RAM) 96
Nways
BroadBand Switch Manager for AIX 566
Campus Manager ATM for AIX 566
Campus Manager for AIX 567
Campus Manager LAN for AIX 566
Manager for Windows 576
Nways Multiprotocol Routing Network Services (MRNS) 89, 98

O

Object Technology (OT) 555
object-oriented technologies 555
Objects comparitions on Java 263
Obscenity 344
ODI 22
offerings 155
Official site security 340
Online directory assistance 558
OPCON 96
Open Graphics Library 229
Open Shortest Path First (OSPF) 103, 131
OpenNet
Americas OpenNet 445
Asian Pacific OpenNet 441
EMEA OpenNet 442
U.S. OpenNet Topology 447
operating system 94
Operating systems 341
operational 520
operational code 92
operational parameters 112
Operations management 387
operator panel display 127
Operators and their precedence. 239
OPTION 189
ordered list 180
Organization Issues 340
organizational impacts 556
originating default 110
OS/2 162
OS/2 drivers 22
OS/2 Platform
DCAF 412
file server 411
IBM SystemView Manager for OS/2 412
LAN NetView Management Utilities 412
LAN Network Manager for OS/2 412
LAN Workgroup 411
main product 412
NetView for OS/2 412
Network Door/2 412
overview 412
positioning 410
skill requirement 411
System Performance Monitor/2 412
OS/2 Warp 159, 160
OS/400 161
OSI 389
OSPF (Open Shortest Path First) 103, 131
OSPF implementation 107
OSPF interoperability 108
outbound filters 136
Overlooked 341

P

Package definition in Java 266
packet fragmentation 45
Panasonic 18
Paper 341
Parallel servers 2
parent node 463
partially meshed 108
partners 130, 149
PASSWORD variable 301
passwords creation 521
PATH_INFO 218
PATH_TRANSLATED 218
PC (Persistent Connections) 44
PDF (Portable Document Format) 205
peer-capable adapters 124
performance 124, 128
Performance management 387
permit mode 137
Persistent Connections (PC) 44
physical network 467
Physical security 348
piggybacking updates 44
pin reset switch 75
PING 103
platform 175
Playing movie files 286
point-to-point (PPP) 94, 100, 102, 130
Point-to-Point Protocol (PPP) 34, 41
pointer resource record 469
Policies 340
Policy 363
POP servers 423, 427

Port 358
Possible problems 346
 POST (power-on self-test) 92
 POST method 216, 224
 Post Office Protocol 3 (POP3) 424
postquery.c 224
PostScript 205
POWER 8
 power switching 43
 power-on self-test (POST) 92
POWER2 8
PowerPC 3, 8
 PPP (Point-to-Point Protocol) 34, 41
 PPP (point-to-point) 94, 100, 102, 130
 PPP compression 32
 Predictor 100
prevent potential system 558
primary link 120
primary name server 469, 472, 481, 482
priority classes 118
Privileges 342
Problem management 387
Procedures 348
project management 557
Proper use 342
Protect and proceed 345
 protocol suite 129
Proxy 355
 Proxy servers 424
PRS 19
PTR resource record 469, 476
 public switched telephone network 21
Pursue and prosecute 345

Q

QAM 18
QPSK 17
 QUERY_STRING 216
query.c 222
QuickTime 287

R

RACF 11
 rack mount options 126
Radio buttons 268
Radio limitations 13
Radio transmission 12
 Rapid Transport Protocol (RTP) 153
RAW audio format 283
rdata resource record 466, 479
 README 99
RealAudio 358
reassembly size 105
Recommendations 4
recursive mode 473, 477
recursive query 473
 redesign 554
REDIRECT_REQUEST 229
REDIRECT_STATUS 229
 reengineering 553
Reflector 15
registering a domain 480
registration 481
 related activities 554
relay hosts 483
 reload 96
 REM (Ring Error Monitor) 101, 145
remapped 511
remapping 510
 remote bridge 132
remote data link switching 116, 149
remote modem 28
 remote server 27
REMOTE_ADDR 218
REMOTE_IDENT 219
REMOTE_USER 219
report option invoking DB2WWW 302
reports 516
 REQUEST_METHOD 215
Requirements 165, 166, 168, 173
resolver general lookup function 476
resolvers 461, 475
resource record definitions 465
resource record query 477
resource records 460, 465
Responsabilities 342
Response 362
 restart 96
restriction check 505
Restrictions in applets 236
REXX 220
RFC 560
RFC 1466 459
RFC 810 460
RFC 822 479
RFC 952 460
RFC, how to obtain a copy 560
RFC, regularly summarizing 560
 RIF (Routing Information Field) 133
 Ring Error Monitor (REM) 101, 145
 Ring Parameter Server (RPS) 79, 101, 145
RIP (routing information protocol) 80, 82, 101, 131
 implementation 106
 interoperability 106
 updates 106
RIPE NCC 480
RISC 1
RISC/6000 163
Risk 341
rlogin 469
 roaming dial-back 28
RODM 568
ROOT 351

root directory 498
root name servers 472
Router and Bridge Manager/6000 565
router cache size 104
router ID 104, 107
Routers and Gateways 425
Routing Information Field (RIF) 133
Routing Information Protocol (RIP) 80, 82, 101, 131
routing table 41, 131
Routing Table Maintenance Protocol (RTMP) 83
routing table size 104
ROW_NUM Variable 306
RPS (Ring Parameter Server) 79, 101, 145
RPT_MAX_ROWS variable 301
RR 465
RS449 93
RTMP (Routing Table Maintenance Protocol) 83
RTP (Rapid Transport Protocol) 153
RTS/CTS 31
Rules 340
run() method 253

S

S-HTTP 159, 160, 370
S/390 1
S/390 security 11
SAP (Service Advertising Protocol) 82
Scalability 3
SCO UNIX 7
Screeaning router 364
Screened host 364
Screened subnet 364
SCRIPT_NAME 218
Scrollbar object 269
SCSI port 127
Second 340
secondary name server 472, 481, 482
Secure browsers 373
Secure Eletronic Payment Protocol 377
secure mail server 483, 484
secure network 482, 483, 484
Secure servers 6, 373
Secure WWW Servers 369
SecurID ACE/Server 86
Security 558
firewall mail server 484
firewall name server 483
firewalls 482
firewalls and electronic mail 483
secure mail server 483, 484
secure network 482, 483, 484
security and DNS 482
security and DNS 482
Security Dynamics ACE/Server 87
Security Dynamics ACE/Server (SecurID) 42
Security incidents 340
Security Mailing Lists 382
segment number 133, 136
SELECT 189
Self-describing audio format 283
SEPP 377
Sequenced Packet Exchange (SPX) 82
Serial Line Internet Protocol (SLIP) 34
Server software 156, 159
SERVER_NAME 218
SERVER_PROTOCOL 218
SERVER_SOFTWARE 218
Service Advertising Protocol (SAP) 82
service port 93
Service Provider 480, 481, 482
Backbone 419, 420
Corporate LAN 423
Corporate Users 422
Customers 422
Firewall 423
Help Desk 421
IBM 8235 DIALs 425
Infrastructure 420, 422
Name Server 424
Network 419
Network Link Speeds 420
Network Solution Design 425
Network Topology 419
Point of Failure 420
Price/Benefit 422
Service Provider 419
Services 422
TCP/IP Network 422
Technical Staff 421
Technology 420
service providers 47
SGML 171
sgmp 389
shared dial-out access 41
SHOWSQL variable 301
showStatus() 276
sibling nodes 463
simple network management protocol (SNMP) 94, 102, 112
simple password 107
Simultaneous users 2
Single Line ISDN Service 19
single-route broadcast 133
Six2View 570
Sizing a server 1
Slip 10, 427
SLIP (Serial Line Internet Protocol) 34
SLS 19
Smart Cards 368
SMP applications 4
SMP Systems 3
SNA 124
data link switching 130, 147
data link switching (DLSw) 100
exchange ID 115

SNA *(continued)*
NetBIOS ratio 152
SNA Manager/6000 565
SNAP (subnetwork access protocol) 136
SNG 373
SNMP 88, 385, 388
Architecture 391
definitions 390
GET operation 400
GET RESPONSE operation 400
GETNEXT operation 400
goals of 392
heterogeneous network 391
history 388
Instance 391
MIB 390
MIB Object 390
OBJECT IDENTIFIER 391
Operations 399
Request/Response Protocol 391
RFC 391
SET operation 400
SNMP Agent 390
SNMP Community 391
SNMP Manager 390
SNMP Model 392
SNMP Proxy Agent 390
SNMP Subagent 390
snmpAuthMsg 398
snmpMgtCom 398
snmpPrivMsg 398
specifications 389
Trap 391
TRAP operation 400
UDP 393
Version 2 397
SNMP (simple network management protocol) 94, 102, 112
SNMP Agent 395
SNMP management 45
SNMP Manager 396
SNMP Subagent 395
SNMP Version 2 397
SOA record 466
SOA resource record 469, 470, 471
Socks 356
Socks servers 424
Software 155, 456
Software bugs 347
software flow control 31
Solaris 7, 163
solution design 457
Somain Name System 459
Sound in Java 263
Source Service Access Point (SSAP) 136
source-route - translational bridging (SR-TB) 100
source-route bridging (SRB) 100, 132
source-route transparent bridging (SRT) 100, 138
spanning tree algorithm 145
Special privileges 344
Splicing 363
spoke and hub network 107
spokes 108, 109
spoofing 44, 114, 149, 363
SPX (Sequenced Packet Exchange) 82
SQL Reports 305
SQL_CODE Variable 306
SR-TB (source-route - translational bridging) 100
SRB (source-route bridging) 100, 132
SRT (source-route transparent bridging) 100, 138
SSAP (Source Service Access Point) 136
SSL 159, 160, 371
Stac 4.0 compression 44
Stacker-LZS 101
Stand-alone 340
start-of-authority (SOA) record 466, 469
static route 110
Static web content 2
Step for Web Development 175
Stereo sound 19
Storage formats 286
String comparison in Java 262
String object 237
String operators 241
stub resolver 477
style 105
subdomain 462, 465, 467, 472, 481, 482
subnetwork 105
subnetwork access protocol (SNAP) 136
subnetwork mask 105
subsystem 97
Superdistribution 527
switch()... case... default 245
Sybase Web.sql 309
symbolic host name 460, 461
symbols 198
sync/async module 76
Synthesizer 285
SYSLOG 349
System managers 341
System Performance Monitor/2 (SPM/2) 574
Systems Monitor for AIX 567
Systemview 385
Application Dimension 387
Data Dimension 387
End-Use Dimension 386
structure and application dimensions 386
structure of 385

T

T1 1, 17, 423
T3 1, 17
table 173, 195
tables 195

TACACS (Terminal Access Controller Access Control System) 64
TAGs 119, 172
Tapes 352
target property 266
TB (transparent bridging) 100, 138
TCP (Transmission Control Protocol) 80, 99, 114
TCP transport layer protocol 146
TCP/IP 94, 155, 432
TCP/IP application 157
TCP/IP client/server 155
Telephone lines 10
Teleprocessing Network Simulator 573
Telnet 80, 93, 346, 357, 476, 557
terminal 93
Terminal Access Controller Access Control System (TACACS) 64
TEXTAREA 190
TextArea object 269
TextField object 269
Textual data 2
TFTP (Trivial File Transfer Protocol) 80, 94, 96, 104, 113
TFTP client 113
TFTP server 113
the Internet 460, 464, 479
Thick (10Base5) 73
Thin (10Base2) 73
third-party security device 28
Threads in java 253
time-to-live 465
time-to-live value 473
Timed LAN-to-LAN Connections (TLC) 44
timed updates 44
Tivoli Management Framework 415
Tivoli TME 10 415
TLC (Timed LAN-to-LAN Connections) 44
TLD 472, 480, 481
TME 10 platforms 416
TMN Workbench for AIX 567
Token-Ring 2, 459
database 141
network bridge program 134
Top-Level Domain (TLD) 464, 480
top-level domains 472
Toshiba 18
toString() method 243
traffic class 118
traffic prioritization 152
Traitors 354
Transformation Services
Business Transformation 553
I/T Strategy 553
Management Technologies 553
Workflow Management 554
translational bridging 125, 141
Transmission Control Protocol (TCP) 80, 99, 114
Transmitting video 279
transparent bridging (TB) 100, 138
Transparent images 205
traps 113
tree 463
triggered updates 44
Trivial File Transfer Protocol (TFTP) 80, 94, 96, 104, 113
Trojan horse 364
Trouble Ticket for AIX 567
try 255
ttl resource record 465
tunnel bridge 122
Tunneling router 364
TUNNELING-IP 119
TXT resource record 466, 469
type resource record 466
Types of MIDI synthesizers 285

U

U.S.Robotics 429
UDP 358
UDP (User Datagram Protocol) 80, 104
Unauthorized access 342
Uniform Resource Locator 489
Uniprocessors 3
University 340
UNIX 3, 463
UNIX Shell Accounts 558
unordered list 181
update() method 258
updates intervals 106
upload files 515
Upstream 17
URL 182, 489
URL prefix 511
URLs on Java 273
USENET 383, 557
User Datagram Protocol (UDP) 80, 104
User responsabilities 344
Users 341
Using existing systems as web servers 2
util.c 220
UTP (10Base-T) 73

V

V_olumn-name 306
V1...Vi Variables 306
variables 113
VC (virtual connections) 44
video 207
Video compression 279
video formats 288
Video quality 279
view 114
VINES (Banyan Virtual Networking Systems) 129

Violated 342
Violated policy 345
virtual connection 54
virtual connections 32, 42, 44, 45
Virtual House Builder 232
Virtual Loadable Modules (VLMs) 41
Virtual network 364
Virtual ROM (VROM) 74
virtual segment number 115, 152
Virus 364, 383
Visual Basic Script Language 233
VLIST Variable 306
VLMs (Virtual Loadable Modules) 41
VRML tutorial 231
VROM (Virtual ROM) 74
VxD 42

W

WAN (Wide Area Network)
Re-Route 101
Restoral (WRS) 101, 119
warm boot 75
Warp Connect 159
Warp Server 4
Warpspace 230
WaveRunner 42, 43
Web applications 556
Web content 508, 521
web farms 452
Web Publisher 163
Web server management software 519
Web server report software 516
web server software 484, 485
Web servers security 520
Web services 557
Web site 169
WebConnection 161
WebExplorer 158
WebTrends 517
WebWATCHER 519
When to Use MIDI 286
while() 247
Windows NT 162
Windows on Java 275
Windows Platform
current product releases 414
key product 414
LAN Remote Monitor for Windows 414
LAN Workgroups 413
NetFinity Manager for Windows 414
NetView for Windows 414
Nways Manager for Windows 414
overview 414
positioning 413
Wireless 12
Word Pro 170, 172, 173
word processors 171
Work Station Gateway
See 5270 to HTML gateway
world wide web 451, 557
World Wide Web Tools 156
WRS (WAN Restoral) 101, 119
www 451

X

X-Windows 358
X.25 94
Xerox Network Systems (XNS) 129
XNS (Xerox Network Systems) 129
XON/XOFF 31

Z

ZModem 94, 95
zone 471, 472, 482
zone filtering 41
zone transfer 473, 482



Printed in U.S.A.

